

Artificial Intelligence Laboratory  
Massachusetts Institute of Technology

AI Memo 389

Logo Memo 37

*The Computer as Coach:  
An Athletic Paradigm for Intellectual Education<sup>1</sup>*

Ira Goldstein<sup>2</sup>

December, 1976

Abstract

Over the next five years, computer games will find their way into a vast number of American homes, creating a unique educational opportunity: the development of "computer coaches" for the serious intellectual skills required by some of these games. From the player's perspective, the coach will provide advice regarding strategy and tactics for better play. But, from the perspective of the coach, the request for help is an opportunity to tutor basic mathematical, scientific or other kinds of knowledge that the game exercises.

Establishing an "athletic" paradigm for skills usually considered the antithesis of ordinary sports is an exciting prospect. There are, however, critical research issues which must be addressed. While the hardware needed for games and coaches will continue to drop in cost, the software technology (and related educational and psychological theory) for designing competent coaches does not yet exist. This is a proposal to develop the theory and design for such coaches, to implement prototypes, and to experiment with their ability to convey important intellectual skills.

<sup>1</sup> This paper is a preliminary proposal submitted to the Science Education Directorate of the National Science Foundation.

<sup>2</sup> Department of Electrical Engineering / Computer Science and the Division for Study and Research in Education.

Contents

<b>Preface</b>	<b>3</b>
<b>1. Design for a Computer Coach</b>	<b>7</b>
1.1 Block Diagram	7
1.2 Theoretical Goals: Towards a Theory of Coaching	9
1.3 Experimental Goals: A Coach for an Elementary Probability Game	11
1.4 A Hypothetical Scenario with a Computer Coach	14
<b>2. Modules of a Computer Coach</b>	<b>22</b>
2.1 The Expert	22
2.2 The Syllabus	23
2.3 The Player Knowledge Model	29
2.4 The Psychologist	29
2.5 The Tutor	32
2.6 The Player Learning Model	40
2.7 The Listener	41
2.8 The Speaker	43
2.9 Summary	44
<b>3. Relevant Research</b>	<b>48</b>
3.1 Computer Assisted Instruction	48
3.2 Cognitive Psychology	52
3.3 Artificial Intelligence	53
3.4 Computer Science	54
<b>4. A Two Year Research Program</b>	<b>55</b>
4.1 Phase I: A Computer Coach for Wumpus	55
4.2 Global Experiments	56
4.3 Subjects	58
4.4 Local AI Experiments	58
4.5 Local Psychological Experiments	59
4.6 Local Pedagogical Experiments	60
4.7 Phase II: Experiments in Other Domains	61
4.8 Phase III: A General Theory of Computer Coaching	62
<b>5. Resources</b>	<b>63</b>
5.1 The MIT Artificial Intelligence Laboratory	63
5.2 The MIT Division for Study and Research in Education	63
5.3 Technology Transfer and Lisp Machines	63
5.4 The Logo Project	64
<b>6. Critique</b>	<b>66</b>
<b>7. Conclusions</b>	<b>68</b>
<b>8. Bibliography</b>	<b>70</b>

Preface

Today TV games for simulated ping pong, soccer and handball are among the most exciting innovations for the home entertainment market. More exciting games in which the player pilots a space ship or drives a race car exist for the commercial market -- recreation rooms in airports, and other public places. These are only the first stones in the approaching avalanche of TV games. Already there are devices that take cassettes on which new programmed games are provided. Among the vast array of computer video games yet to be marketed, some will exercise serious mathematical and scientific knowledge.

In PING PONG, SOCCER, and HANDBALL, the player controls a paddle so as to hit a moving "ball" into an opponent's goal or other such target. These first generation games provide a limited intellectual environment. Second generation games in which the player controls a space craft, boat or a race car are still too expensive for the home. These games involve controlling a vehicle given the complication of skidding for cars, drifting for boats and falling for space ships. They make greater intellectual demands on the player. Successful navigation requires knowledge of geometry, dynamics and kinematics. Third generation games will have available powerful computational resources: the possibilities are myriad. For example, consider STEVEDORE, a hypothetical game that illustrates the range of possibilities. In this game, the player is asked to load a cargo, given various sets of simple machines. The machines have costs associated with them: the task is to find the cheapest combination of simple machines adequate to move the weight to the desired location. Successful play involves in a natural and active way knowledge of elementary physics. Finally, a set of third generation computer games already exists, having been developed in the context of efforts like PLATO (such as HOW THE WEST WAS WON). These can be translated to the home market.

Such games have a powerful educational appeal. They will be:

1. widespread: the calculator phenomenon of drastically reduced prices is about to be repeated for computer-based TV games. Research in education can try to take advantage of this phenomenon, or ignore it; but it will happen in any case.
2. active: knowledge learned is used for a purpose. What angle of the paddle will establish the desired trajectory of the ball; what force is necessary to enter a stable orbit; what combination of simple machines can lift the desired weight? The passive environment of the traditional classroom or educational television is avoided.
3. motivating: computer games will be played because they are enjoyable, not because of some external demand made on the student. The desire for instruction to improve his play arises naturally on the part of the player.

However, games have a limitation: the player, on his own, can fail to acquire the skills of an expert. This suggests that coaches be developed for computer games. From the player's perspective, the coach will provide advice on strategy and tactics for better play. But, from the perspective of the coach, the request for help is an opportunity to tutor basic mathematical, scientific or other kinds of knowledge that the game exercises.

Human coaches are possible, except that the games will be so widespread that it will be difficult to supply the required number of skilled teachers. Furthermore, the games are often dynamic, making it difficult for a human coach to follow the play in real time. Hence, our proposal is to develop and test computer coaches.

There is another virtue to the design of a computer coach -- the rigor required to write a program provides a controlled environment to study basic



questions of learning and teaching. Insight into these questions will have theoretical value for education extending beyond the direct application of computers.

Thus, we propose to develop the theory and design for such coaches, to implement prototypes and to experiment with their ability to convey important intellectual skills. Specifically, we will address these questions:

1. How can we design expertise in the coach so that it can respond reasonably to the player's particular choice of move? If the coach forced the player to view the game in only one way, it would be a straight-jacket. To avoid this, the coach must be able to analyze a wide variety of moves and discuss their relative merits. Progress in artificial intelligence (AI) makes this a possible goal for the closed world of a game. We propose to apply AI to the theory and design of an "Expert" component in the coach.
2. How can we model the player sufficiently well so that the coach's remarks are appropriate, neither too advanced for a beginner nor too elementary for an expert? Here we propose to apply the formal modelling tools of information processing psychology to the theory and design of a "Psychologist" component in the coach.
3. How can we control the nature of the coach's advice so that it is given in a friendly and personal manner? The solution lies in having: an array of possible interaction modes ranging from graphics to natural language; a theory of how to abbreviate complex explanations; and a model of the player's learning preferences. Both AI and information processing psychology will be used to design a "Tutor" component of the coach with these capabilities.

The design of a successful computer coach is a difficult enterprise. But we are about to experience the explosive diffusion of computer game technology. If

we are successful in taking educational advantage of this, the rewards will be enormous.

Chapter 1 outlines our design for a Computer Coach in terms of modules that have responsibility for domain expertise, for modelling, for tutoring and for generating English prose. Each module is based upon a rule-based formulation of the appropriate knowledge. Chapter 2 provides details.

Chapter 3 reviews relevant research in AI, information processing psychology, and computer aided instruction. The computer coach owes its greatest intellectual debt to the work of J. S. Brown and his colleagues who have pioneered the design of computer-based tutors for various domains.

Chapter 4 describes experiments to evaluate the computer coach paradigm as a vehicle for tutoring transferable intellectual skills. Phase I involves the implementation and testing of a coach for an elementary probability game. Phase II analyzes the paradigm for other games. Phase III is addressed to the articulation of a general procedural theory of coaching. Since phases II and III are dependent on the success of phase I, support is requested only for phase I as a two year research project. At the end of this period, the progress made on phase I will determine the appropriateness of a new proposal requesting support for phases II and III.

## 1. Design for a Computer Coach

### 1.1 Block Diagram

Figure 1 is a block diagram of the design for a computer coach.<sup>1</sup> We have used anthropomorphic designations for the Expert, Tutor and Psychologist modules of the Coach to emphasize their purpose in the overall system. Of course, these components will be far more limited than their human counterparts, although we believe it is possible to get significant performance. To stress that we are referring to computer programs and not people, we capitalize references to these modules in the text.

The function of the Coach is to tutor the player in particular skills, in the context of situations where those skills are applicable. It does this through the interactions of the specialist programs appearing in the block diagram. The Expert informs the Tutor if the player's move is nonoptimal and which skills are needed to discover better alternatives. These skills are potential topics for the Tutor to discuss. The Psychologist examines the student's behavior and makes hypotheses about which skills are already possessed by him (recorded as the Knowledge model) and which tutorial modes are effective in conveying new skills to him (recorded as the Learning model). The Tutor uses these models to personalize its interactions with the player. The Knowledge model guides the selection of topic from those suggested by the Expert while the Learning model influences the choice of explanatory strategy. Finally, a Speaker converts the formal explanation of the Tutor to an appropriate form for the player. Usually this would be English, though graphical explanations are also possible using the TV as a display.

<sup>1</sup> This design has grown out of a close collaboration with B. Carr, M. Miller and J. Stansfield of MIT and J. Brown and A. Collins of Bolt, Baranek and Neuman.

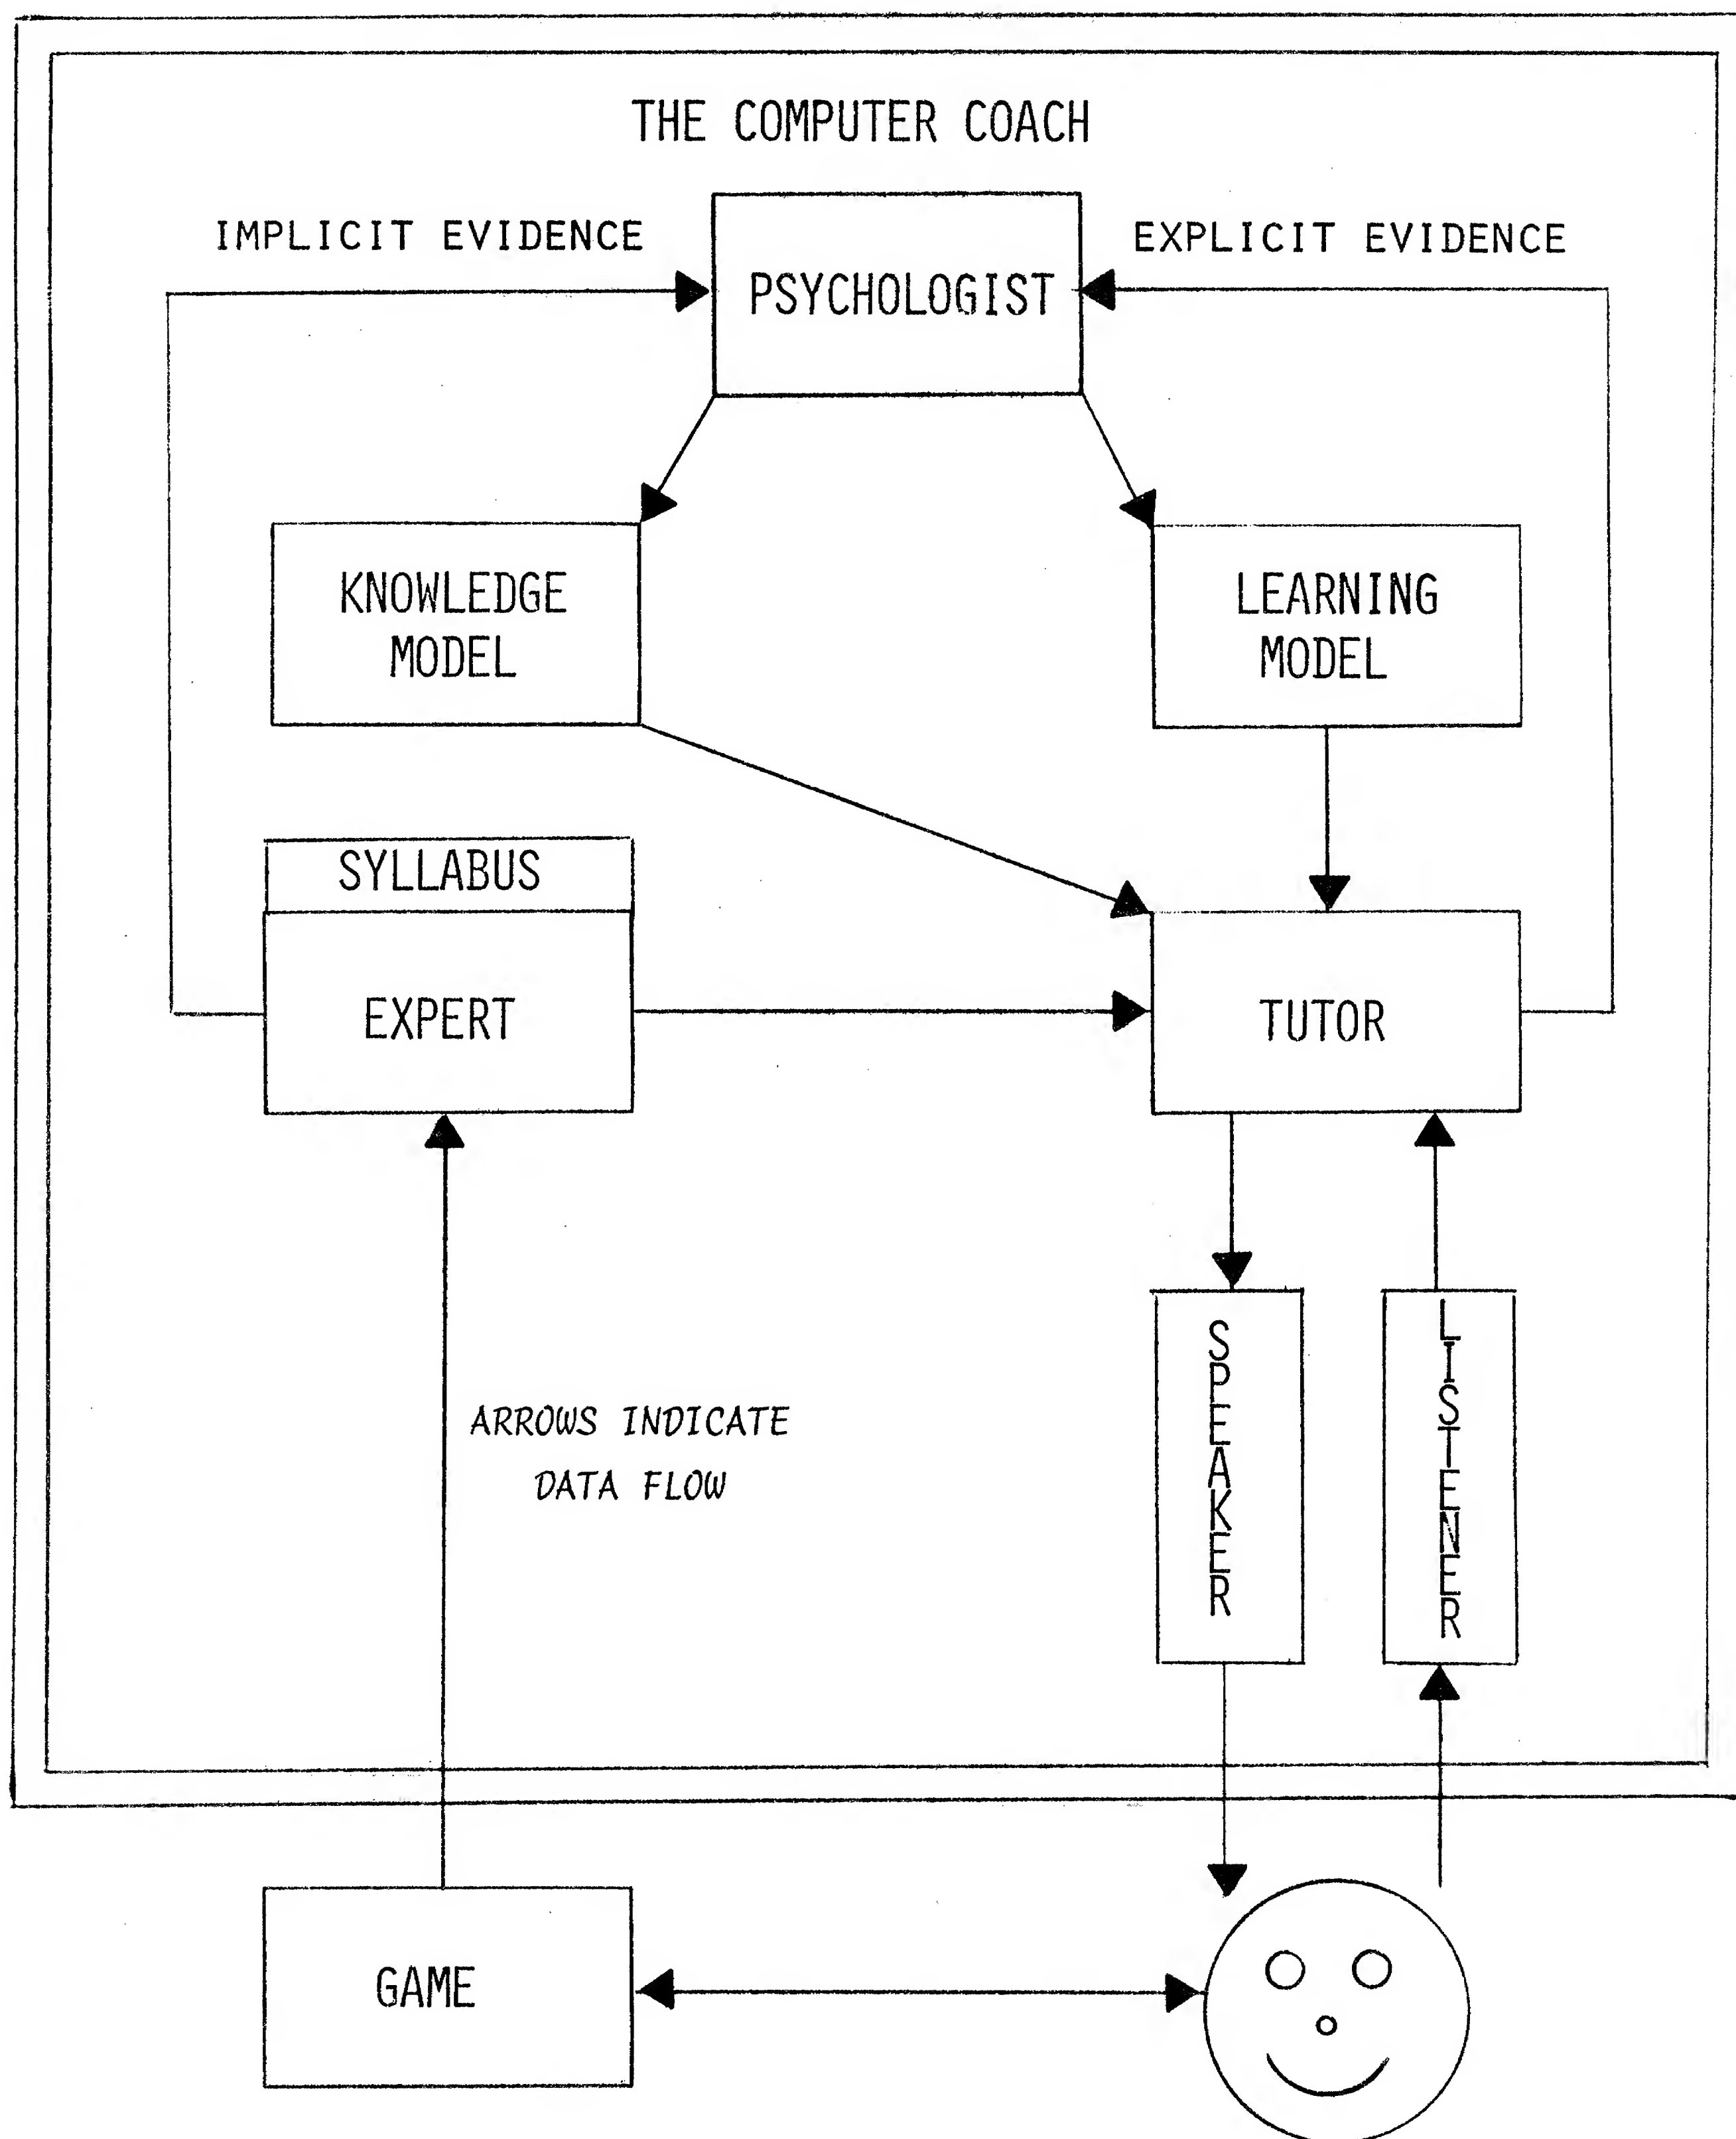


FIGURE 1



## 1.2 Theoretical Goals: Towards a Theory of Coaching

The theoretical goal of this research is to study fundamental questions in the theory of learning, modelling and teaching by constructing procedural rule systems for (1) the skills needed by the Expert to play the game, (2) the modelling criteria used by the Psychologist, (3) the alternative tutoring strategies used by the Tutor, and (4) the language generation capabilities of the Speaker. We expand on each below:

1. The Expert will use rules of skill which embody the knowledge required to play the game and thereby analyze the player's behavior. The virtue of a rule-based representation of expertise is that its modularity allows tutoring to be focussed concisely on the discussion of specific skills, and permits modelling to take the form of hypotheses regarding which rules are known by the player.

A possible confusion should be clarified here. When we refer to the rules of the Expert, we are not referring to the "rules of the game", i.e. the facts describing a legal move and what constitutes a winning state. Rather our concern is with the tactical and strategic knowledge needed to decide which move among the legal possibilities to make. We term these "rules" because our representation methodology is to structure the skills in terms of rule sets.

2. The Psychologist will use rules of evidence to make reasonable hypotheses about which skills of the Expert the player possesses. Typical rules would be:

A. Increase the estimate that a player possesses a skill if the player explicitly claims acquaintance with the skill; and decrease the reliability if the player expresses unfamiliarity.

B. Increase the estimate that a player possesses a skill if the skill is manifest in the player's behavior; decrease the estimate if the skill is not manifest in a situation where the Expert believes it to be appropriate.

Hence, implicit as well as overt evidence plays a role.

- C. Decrease the estimate that a player possesses a skill if there is a long interval since the last confirmation was obtained (thereby modelling the tendency for a skill to decay with little use).
3. The Tutor will use rules of explanation to select the appropriate topic to discuss with the player and to choose the form of the explanation. These rules include:
- A. Rules of simplification that take a complex statement and reduce it to a simpler assertion. Simplification rules are essential if the player is not to be overwhelmed by the Tutor's explanations.
  - B. Rules of rhetoric that codify alternative explanation strategies. Two extremes are explanation in terms of a general rule versus explanation in terms of a concrete instance.
4. The Speaker will use rules of language to convert the formal message selected by the Tutor to linguistic form. This involves an AI language generator which we discuss later in the paper. Various mechanisms of language for achieving brevity such as anaphora and ellipsis would be applied.

These research areas are difficult, touching upon deep issues in psychology, education, linguistics and artificial intelligence. But, we believe there is an opportunity for progress for three reasons:

- 1. Modelling and tutoring are being examined in the constrained context of a game. A game has a formal structure, a restricted number of options, and involves a limited number of skills. Furthermore, these constraints make it possible to build a competent Expert for the domain.

2. The research is an integrated enterprise with the potential for a synergistic effect. For example, the Psychologist uses the same simplification rules to generate a simplified version of the Expert for the initial model of the Player, as the Tutor uses to summarize an explanation. Similarly, the overall methodology of representing knowledge as procedural rule sets is being applied to all of the components of the Coach.
3. The rule-based computational paradigm which we intend to exploit is a powerful one and is particularly appropriate to the dynamic demands of a theory of interaction between coach and player.

### 1.3 Experimental Goals: A Coach for an Elementary Probability Game

In view of the many difficult issues raised by our theoretical goals, the experimental focus of this proposal will be restricted to developing a Coach for a single example of an intellectual game -- Wumpus. (Although the tutoring and modelling components of the Coach will be designed in a modular, domain-independent fashion which will be transferable to a wide range of tasks.) A set of experiments described in chapter 4 will test:

1. the relative merits of alternative designs for the modules of the Coach, and
2. the overall success of the Coach in facilitating the acquisition of transferable intellectual skills.

The Wumpus game was created by Gregory Yob [1975], as an improvement to other games he had seen being played by computer hobbyists. (Thus, it is further evidence for the manifold possibilities that exist for the computer game environment.) It exercises basic knowledge of logic, probability, decision analysis and geometry. It is not yet available for the home; but, judging by the enjoyment our MIT implementation of the game (with no coaching) has provided to players ranging in age from elementary school children to adults, we think it



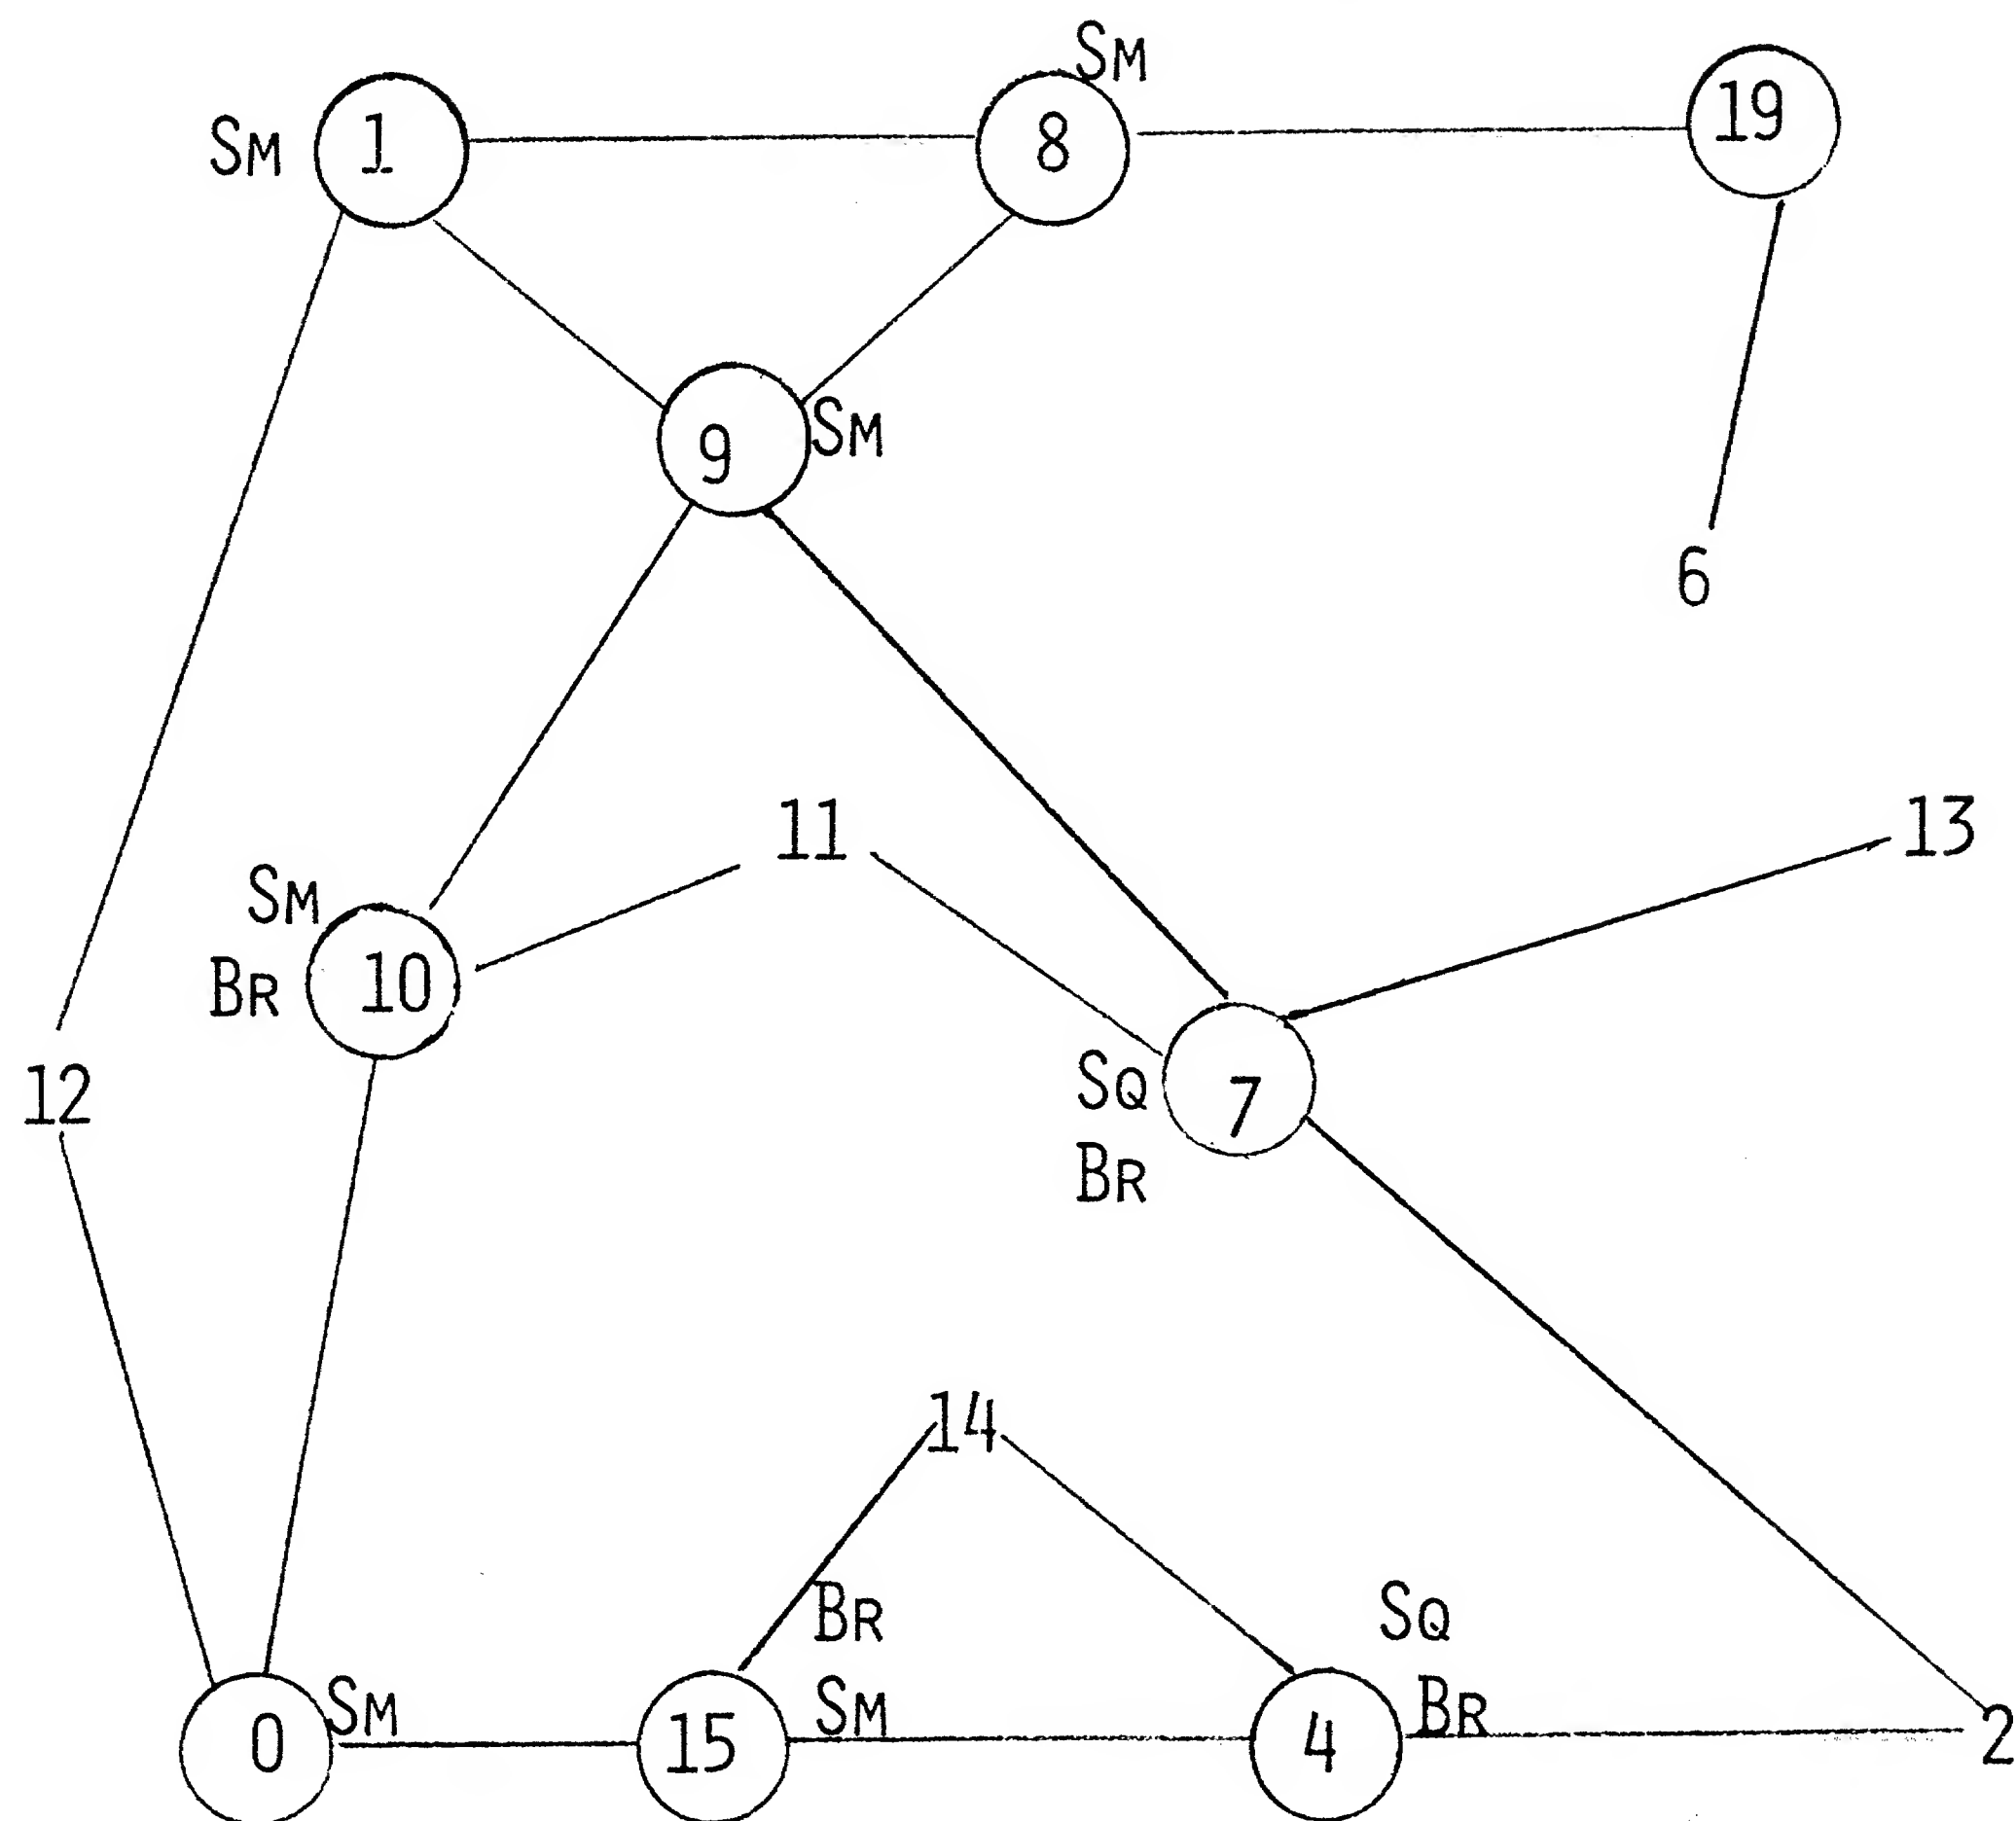
inevitable that activities of this kind will appear among the third generation of computer games. (Its demands on computer resources are slight.)

The game is a modern day version of Theseus and the Minotaur. The player is initially placed somewhere in a randomly connected warren of caves and told the neighbors of his current location. His goal is to locate the horrid Wumpus beast in the warren and slay it with an arrow. Each move to a neighboring cave yields information regarding that cave's neighbors. The difficulty in choosing a move arises from the existence of dangers in the warren -- bats, pits and the Wumpus itself. If the player moves into the Wumpus' lair, he is eaten. If he walks into a pit, he falls to his death. Bats pick the player up and randomly drop him elsewhere in the warren. But the player can minimize risk and locate the Wumpus by making the proper logistic and probabilistic inferences from warnings he is given. These warnings are provided whenever the player is in the vicinity of a danger. The Wumpus can be smelled within one or two caves. The squeak of bats can be heard one cave away and the breeze of a pit felt one cave away. The game is won by the player shooting one of his arrows into the Wumpus' lair from an adjoining cave. If he exhausts his set of five arrows without hitting the creature, he has lost the game. Figure 2 illustrates a typical intermediate state a player might reach.

Skilled play exercises basic skills in:

- A. logic -- making deductions in those situations where complete knowledge is available, e.g. realizing that all neighbors of a cave are safe if no warning is received,
- B. probability -- selecting the best choice given uncertain knowledge, e.g. knowing that the likelihood that a cave contains a danger increases if there are warnings from more than one neighbor.





CIRCLED CAVES HAVE BEEN  
VISITED BY PLAYER.

SM = WUMPUS WARNING  
SQ = BAT WARNING  
BR = PIT WARNING

FIGURE 2  
AN INTERMEDIATE STATE IN A TYPICAL WUMPUS GAME

- C. decision making -- selecting the move with highest utility by balancing information gain against increased danger, and
- D. geometry -- in some variations of the game, deducing constraints from the geometry of the maze is possible. The game can be played in warrens ranging from general 3-dimensional lattices to 2-dimensional rectangular grids.

In addition to being a motivating and intellectually challenging game, Wumpus also has the virtue that we have already had some experience with this game. Our experience derives from a course given by the Division for Study and Research in Education at MIT by the author and J. Stansfield. The class project was to design and implement a simple advisor for the Wumpus game [Stansfield, Carr & Goldstein 1976]. This advisor called WUSOR has been subsequently developed by B. Carr. He has implemented a powerful rule-based expert for the game with limited modelling and tutoring capabilities. The design proposed here is based on that experience; however, it represents a significant extension along the dimensions of improved modelling, tutoring and speaking abilities. The nature of these extensions are the subject of chapter 2. Furthermore, no version of the Wumpus coach has, as yet, been subject to rigorous evaluation. Chapter 4 outlines a thorough experimental program to meet this goal.

#### 1.4 A Hypothetical Scenario with a Computer Coach

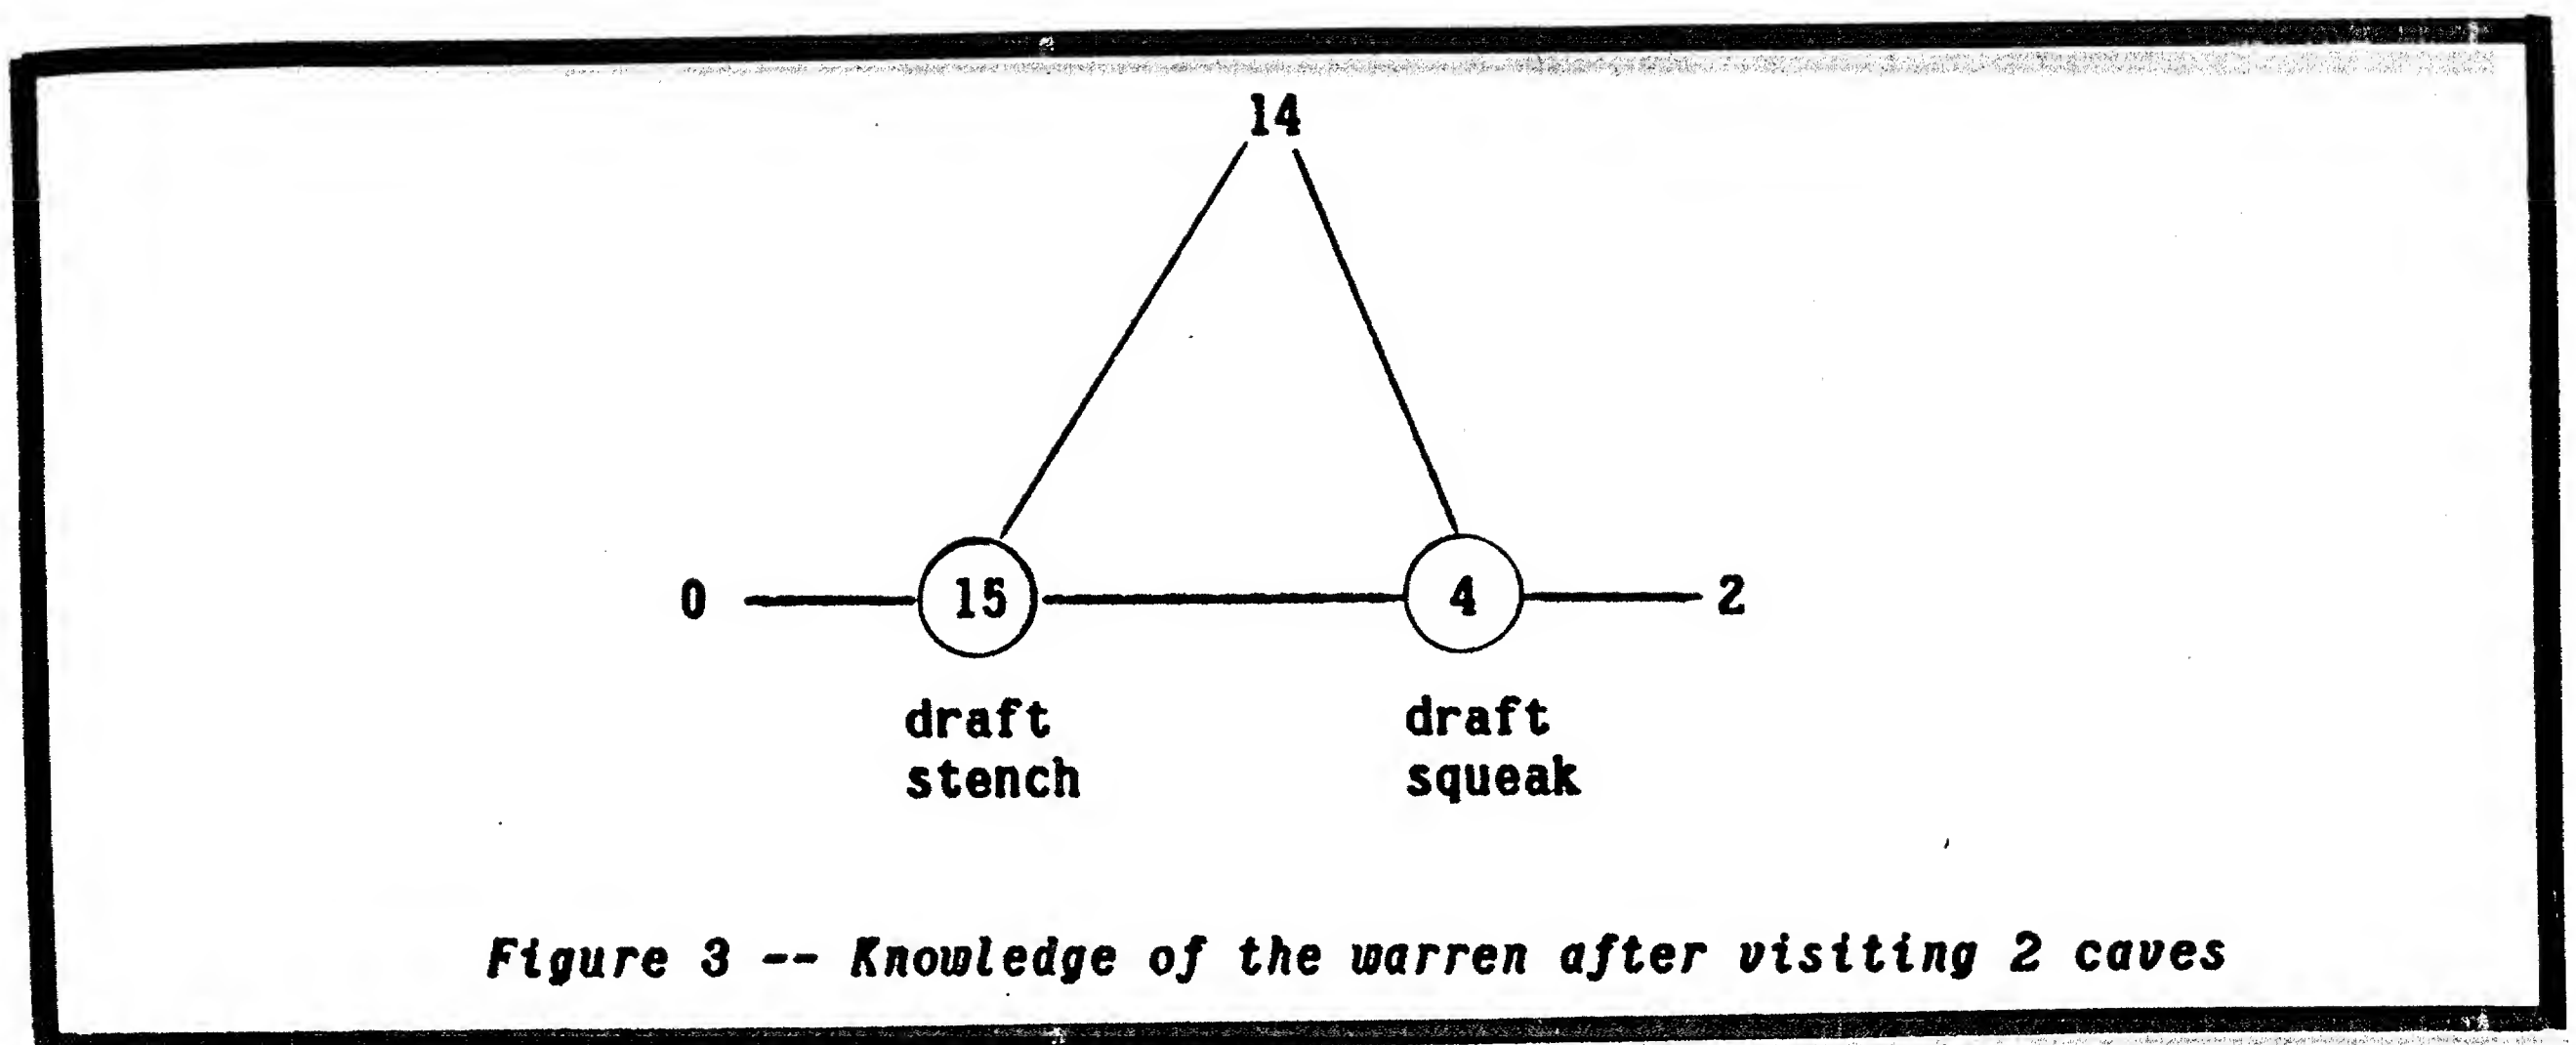
To illustrate the potential use of a coach as well as the intellectual skills involved in playing the Wumpus game, we provide a hypothetical scenario of a player interacting with the proposed Coach (henceforth called COACH-1). The player's responses are preceded by a ">".

*You are now at cave 15 with neighbors 4, 14 and 0. Brrr! There is a draft. You are near a pit. What a stench. The Wumpus is near! What now?*

*> 4*

*You are now at cave 4 with neighbors 15, 14 and 2. Brrr! Squeak! A bat is near. What now?*

Figure 3 is a typical picture drawn by a player to record the information learned about the warren.



The goal of the Coach is to tactfully tutor a beginner in the relevant logical, probabilistic and strategic knowledge needed to play the game. For example, on the basis of the above knowledge of the warren, cave 14 should be treated as more dangerous than 0 or 2 since there is multiple evidence (from the drafts in 15 and 4) that 14 contains a pit. If the player moved to cave 14, a coaching situation might occur as follows:

*> 14*

*Gooooobyeeee! You have fallen in a pit. What now -- coaching  
or a new game?*

*> Coaching*

*Cave 0 was safer.*

*> Why?*

*You had single evidence that it contained a pit, but double  
evidence for cave 14.*

*> Is cave 0 a better move than cave 2.*

*Yes. Cave 2 contains a bat.*

*> New game.*

Although it is not apparent from these simple remarks, every module of the Coach contributed to the dialog. We summarize these interactions here under the headings -- tutoring, modelling, explanation, discourse, and natural language comprehension; then develop the theory and design for each component of the Coach in the next chapter.

1. Tutoring: The Coach had the option of interrupting immediately upon the player's move to cave 14. Instead, it allowed the player to learn directly of his fatal move. The decision whether to intervene immediately, or only upon request, is made by the Tutor on the basis of the Learning model. The Learning model records the Psychologist's hypotheses regarding whether the player prefers immediate intervention or not. In this case, previous experience with the player resulted in the belief that intervention only upon request was desired. The Tutor also considers whether termination of the game should be avoided because other learning opportunities exist in this particular state of the game. This



would be determined on the basis of the Expert's analysis of the skills needed to play the game in its current state as compared with the Knowledge model's hypotheses regarding those skills in need of practice.

2. Modelling: The move to cave 14 causes the Psychologist to decrease the Knowledge model weight indicating familiarity with the Double Evidence rule. (This is the Expert's rule that assigns an increased probability of danger to a cave for which there is double evidence. A sampling of the Expert's inference rules are given in the next chapter; the Knowledge model contains a numerical estimate of the Player's knowledge of each rule.)

Modelling raises many issues. One subtlety is that the move to 14 may be evidence of a more elementary limitation -- a failure to understand the logical implications of the squeak warning, i.e. that a bat is in a neighboring cave. The current state of the Knowledge model is used by the Psychologist to determine which skill is missing when a nonoptimal move is made. The Knowledge model indicates the level of play which can be expected from this player -- the player might be a pure beginner with incomplete knowledge of the basic rules of the game, a novice with understanding of the logical skills, an amateur with knowledge of the logical and the more elementary probability skills, etc. The Psychologist would attribute the unfamiliarity to an skill at the student's current level of play -- in this case, we are presumably dealing with a novice player who has mastered the logical skills and is learning the basic probability heuristics.

Another subtlety arises from potential interactions between the player's choice of representation scheme and his application of the logical and probability skills to the information contained in that representation. He might know the double evidence rule, but have represented the information incompletely, and hence have not recognized its applicability. The Psychologist can choose to control for this by providing a graphic representation for the Player.

3. Explanation: The response to the player's initial request for coaching -- "Cave 0 was safer." is abridged. The reason for this is that the complete explanation may be overwhelming to the player. For example, the Expert will be capable of the following complete analysis (expressed here in English, although the Expert's analyses will be formal derivations, resembling mathematical proofs rather than text):

*Cave 14 was not the best move. Logically, the Wumpus cannot be in 0, 2 or 14 since there is no smell in 4. But caves 0 and 2 were better than 14 because there was single evidence that caves 0 and 2 contained a pit, but double evidence for cave 14. Finally, cave 2 is more dangerous than cave 0, since 2 contains a bat, and the bat could drop you in a fatal cave. (No squeak in 15 rules out the possibility of a bat in 14; hence, the squeak in 4 can only be explained by a bat in 2.) Thus, the best move is to cave 0.*

Giving a complete explanation does not encourage reflection on the player's part. Hence, the Tutor prunes the complete analysis on the basis of simplification rules and provides only a headline. Additional information is given only if requested by the player.

4. Discourse: Further brevity is obtained by the Speaker module's use of ellipsis and anaphora in generating English replies for the player. For example, the Coach begins with the elliptical utterance "Cave 0 is safer", rather than the complete clause "Cave 0 is safer than cave 14." The ellipsis is justified by the context, in which the player has just moved into the fatal cave 14. In the next statement by the Coach, "You had single evidence that it contained a pit, ...", a pronoun is used to refer to cave 0.

The underlying formal explanation of the Tutor would have had this

repetition, as it will be composed of logically complete statements regarding the rules in analyzing the risk of various moves. In the absence of any discourse rules, the Speaker would have mirrored this repetition in the generated English.

The reader may feel we are making too strong a virtue of brevity. However, the ability to make concise, appropriate remarks is a critical capability of the Coach. Otherwise the enjoyment of the game and the efficacy of the tutor will suffer. Rules for simplification, both for the underlying conceptual level and the surface syntactic level, are an important part of a general theory of coaching.

Of course, the Coach can err on the side of too much brevity. Confusion can result from saying too little as well as too much. This is partially alleviated by giving the player the ability to ask questions, which we discuss next. It is also addressed by incorporating into the Learning model the Psychologist's hypotheses regarding whether the explanations are proving satisfactory to the player, as estimated by the player's responses and subsequent behavior. On the basis of these hypotheses, the use of particular simplification rules can be adjusted.

5. Natural Language Comprehension: While most of the player's responses were simple one word remarks, his last response was the question: "Is cave 0 safer than cave 2?" The Listener, using standard AI techniques, parses this sentence into a formal representation that indicates that the student is asking a question about two alternative moves. This raises the question of what class of utterances the Coach can be expected to comprehend.

Our plan is for the system to understand those formal queries that already occur on the communication channels between the Psychologist, Tutor and Expert, i.e. one can conceptualize the Tutor, Psychologist and Expert as constantly asking questions one of another. The student's question is answerable if it

falls among this class.

From this perspective, the job of the Listener is to convert the remark to its formal form, or decide that it is not comprehensible. Winograd's SHRDLU program [1972] and Woods' LUNAR system [1972] can perform such conversion for a reasonable range of English constructions, providing there is a well-defined discourse world. The set of formal queries permissible between modules of the Coach is such a world. Thus, the design of the Coach as a community of communicating specialists is justified both by the virtues of modularity and by the support it provides for a language capability.

In the scenario, the question "Is cave 0 better than cave 2?" requires an analysis of the relative merits of alternative moves -- a capacity basic to the Expert. Thus, the Coach can respond appropriately.

(The ability to be articulate about its own thinking is critical for another reason: if a human teacher is to accept the Coach as an aid, he or she must be convinced of its competence. Requiring that the teacher examine the code is absurd. Instead, we envision the teacher pretending to be a student, and then demanding explanations from the Coach for its behavior. Hence, the design philosophy of communicating specialists is essential if the Coach is to be able to explain itself.)

Finally, a remark about the complexity of the Wumpus game is appropriate. If the player had not moved to cave 14, the game might have continued until an intermediate state such as that shown earlier in figure 2 was reached. At this point, the game is quite challenging; for example, the reader may be surprised to learn that a careful application of the logical rules of the game allows one to deduce that the Wumpus is in cave 12. However, the game does not become significantly more complex for the Coach. The Expert remains able to analyze the



probability of danger for all possible moves to unvisited caves with the bookkeeping abilities of the computer preventing confusion. The tutoring and modelling continues to be focussed on the difference between the player's move and any better moves the Expert discerns.

The next chapter discusses the proposed design for each module of the coach.

Modules of a Computer Coach2.1 The Expert

The block diagram of figure 1 shows a prominent module for the domain expert. The incorporation of such an Expert represents a basic insight of work in recent CAI: intelligent tutoring requires knowledge of the subject matter. The power of such knowledge is seen in tutors for geography (Scholar [Carbonell 1970]), electronics (Sophie [Brown & Burton 1975]), set theory (Excheck [Smith et al, 1975]), and arithmetic (West [Burton & Brown 1976]). At MIT, we have developed new models of expertise for planning and debugging [Goldstein & Grimson, in press; Miller & Goldstein 1976a].

An Expert for Wumpus implemented by Carr will provide the needed understanding of the player's options that must be available to the Coach. The Expert's knowledge consists of a set of 20 rules describing various logical, probabilistic, geometric and strategic facts about the game, a few of which are:

Logical Rules for Bats & Pits

- L1: (positive evidence rule) A warning in a cave implies that a danger exists in a neighbor.
- L2: (negative evidence rule) The absence of a warning implies that no danger exists in any neighbors.
- L3: (elimination rule) If a cave has a warning and all but one of its neighbors are known to be safe, then the danger is in the remaining neighbor.

Probabilistic Rules for Bats & Pits

- P1: (equal likelihood rule) In the absence of other knowledge, all of the neighbors of a cave with a warning are equally likely to contain a danger.
- P2: (double evidence rule) Multiple warnings increase the likelihood that a given

*cave contains a danger.*

The rules are given explicitly for bats and pits, but the same implications are true for the Wumpus, except that warnings propagate two caves. Note that these "rules" are not the rules of the game, but rather inference rules from which the best move can be deduced. The inference rules include the rules of the game plus general strategic knowledge about probabilities and constraints. They are essentially rules of scientific induction.

The experimental goal of this proposal is to develop sophisticated modelling and tutoring capabilities. However, this goal does not diminish the importance of the Expert. Indeed, the *first design principle* for such Coaches is that the nature of the other components of the Coach depends in critical ways on the Expert -- the Psychologist models the player in terms of subsets and simplifications of the Expert's knowledge while the Tutor selects an appropriate issue partly by referencing a Syllabus representing an ordering on the Expert's knowledge.

## 2.2 The Syllabus

A syllabus is needed for the coach to determine which intellectual issues to discuss with the player. In figure 1, the syllabus appears on top of the Expert. This is to emphasize that a rule based theory of expertise allows the development of a syllabus in terms of subsets and simplifications of the Expert's rules.

Selection of subsets of the Expert's knowledge as intermediate goals for the Coach is based upon the complexity of the various rules and their interdependencies. For example, the logical and probabilistic rules of Wumpus form two subsets of the Expert's knowledge, with the natural tutorial sequence being to begin with the logical rules. This is required by the nature of the rules, since knowledge of the logical rules is needed to properly apply the probability heuristics.

The utility of a syllabus was exemplified by the scenario of the previous chapter. The Coach assumed that the nonoptimal move was due to lack of familiarity with the Double Evidence probability rule. However, that rule cannot be properly applied unless the necessary logical deductions regarding the available evidence are made. Therefore, for a player whose experience and/or age places him early in the syllabus, the logical rules would be the preferred topic of conversation. The scenario involved a more advanced player, who had already mastered the logical rules; hence, the coaching focus was on the next topic of the syllabus -- the probability rules.

Creating subsets of the Expert's knowledge is straightforward: deciding on useful simplifications of various rules is more subtle. This is true for the traditional educational setting as well as for a Computer Coach. An example of a rule simplification is as follows: Suppose a rule has certain exceptions. A reasonable pedagogical simplification is for the Coach to ignore these exceptions, until the novice player has mastered the exception-free approximation. For Wumpus, a typical simplification of this kind is to assume that a given warning is caused by only a single danger. The more subtle analysis needed to deal correctly with the existence of multiple dangers of each type (i.e. multiple bats, pits and Wumpii) is for a later stage in the syllabus.

Since simplification is a critical ability, the investigation of general criteria for simplifying a rule-based expert to yield a syllabus becomes an important goal. One candidate is to simplify propositions of the form "A and B", "A or B" or "B => A" as "A". This general schema produces the simplification of assuming that a single danger is responsible for multiple evidence. Work by Rumelhart [1975] on summarization rules for stories suggests that the simplification rules proposed here (which are similar to his rules) have some



psychological validity in terms of how people simplify and remember descriptions.

The Coach will take advantage of simplified subsets of the Expert's rules in four ways. The first is to form "Simulated Players" from these subsets. These would represent average players at various stages of development. Simulated Players would augment the Expert by providing a source of expectations for student-level performance. The Simulated Player would also represent a target knowledge state for players at a preceding stage in the skill acquisition process. Figure 4 illustrates the augmented Computer Coach.

The second use of simplification rules is to derive variations on the Wumpus game as intermediate coaching environments. Such variations can be obtained by changing the number and type of dangers, the distance from which they can be sensed, the number of caves and the topological restrictions on the warren (i.e. a general three dimensional maze at one end of the spectrum to a 2 dimensional rectangular grid at the other). Versions differ in their complexity and the skills they require; hence good coaching includes the capability of suggesting the appropriate variation to a player, depending on his level of skill. For example, the following sequence of successively more complex variations might be used by the coach:

1. (single, static dangers in constrained maze) Wumpus with only a single danger of each type, played on a rectangular grid.
2. (multiple, static dangers in constrained maze) Multiple bats and pits, but only a single Wumpus, again played on a rectangular grid.
3. (multiple static dangers in a general maze) Generalizing the grid to an unconstrained maze with a limit on the branching factor at any cave.

## COMPUTER COACH WITH SYNTHETIC PLAYER

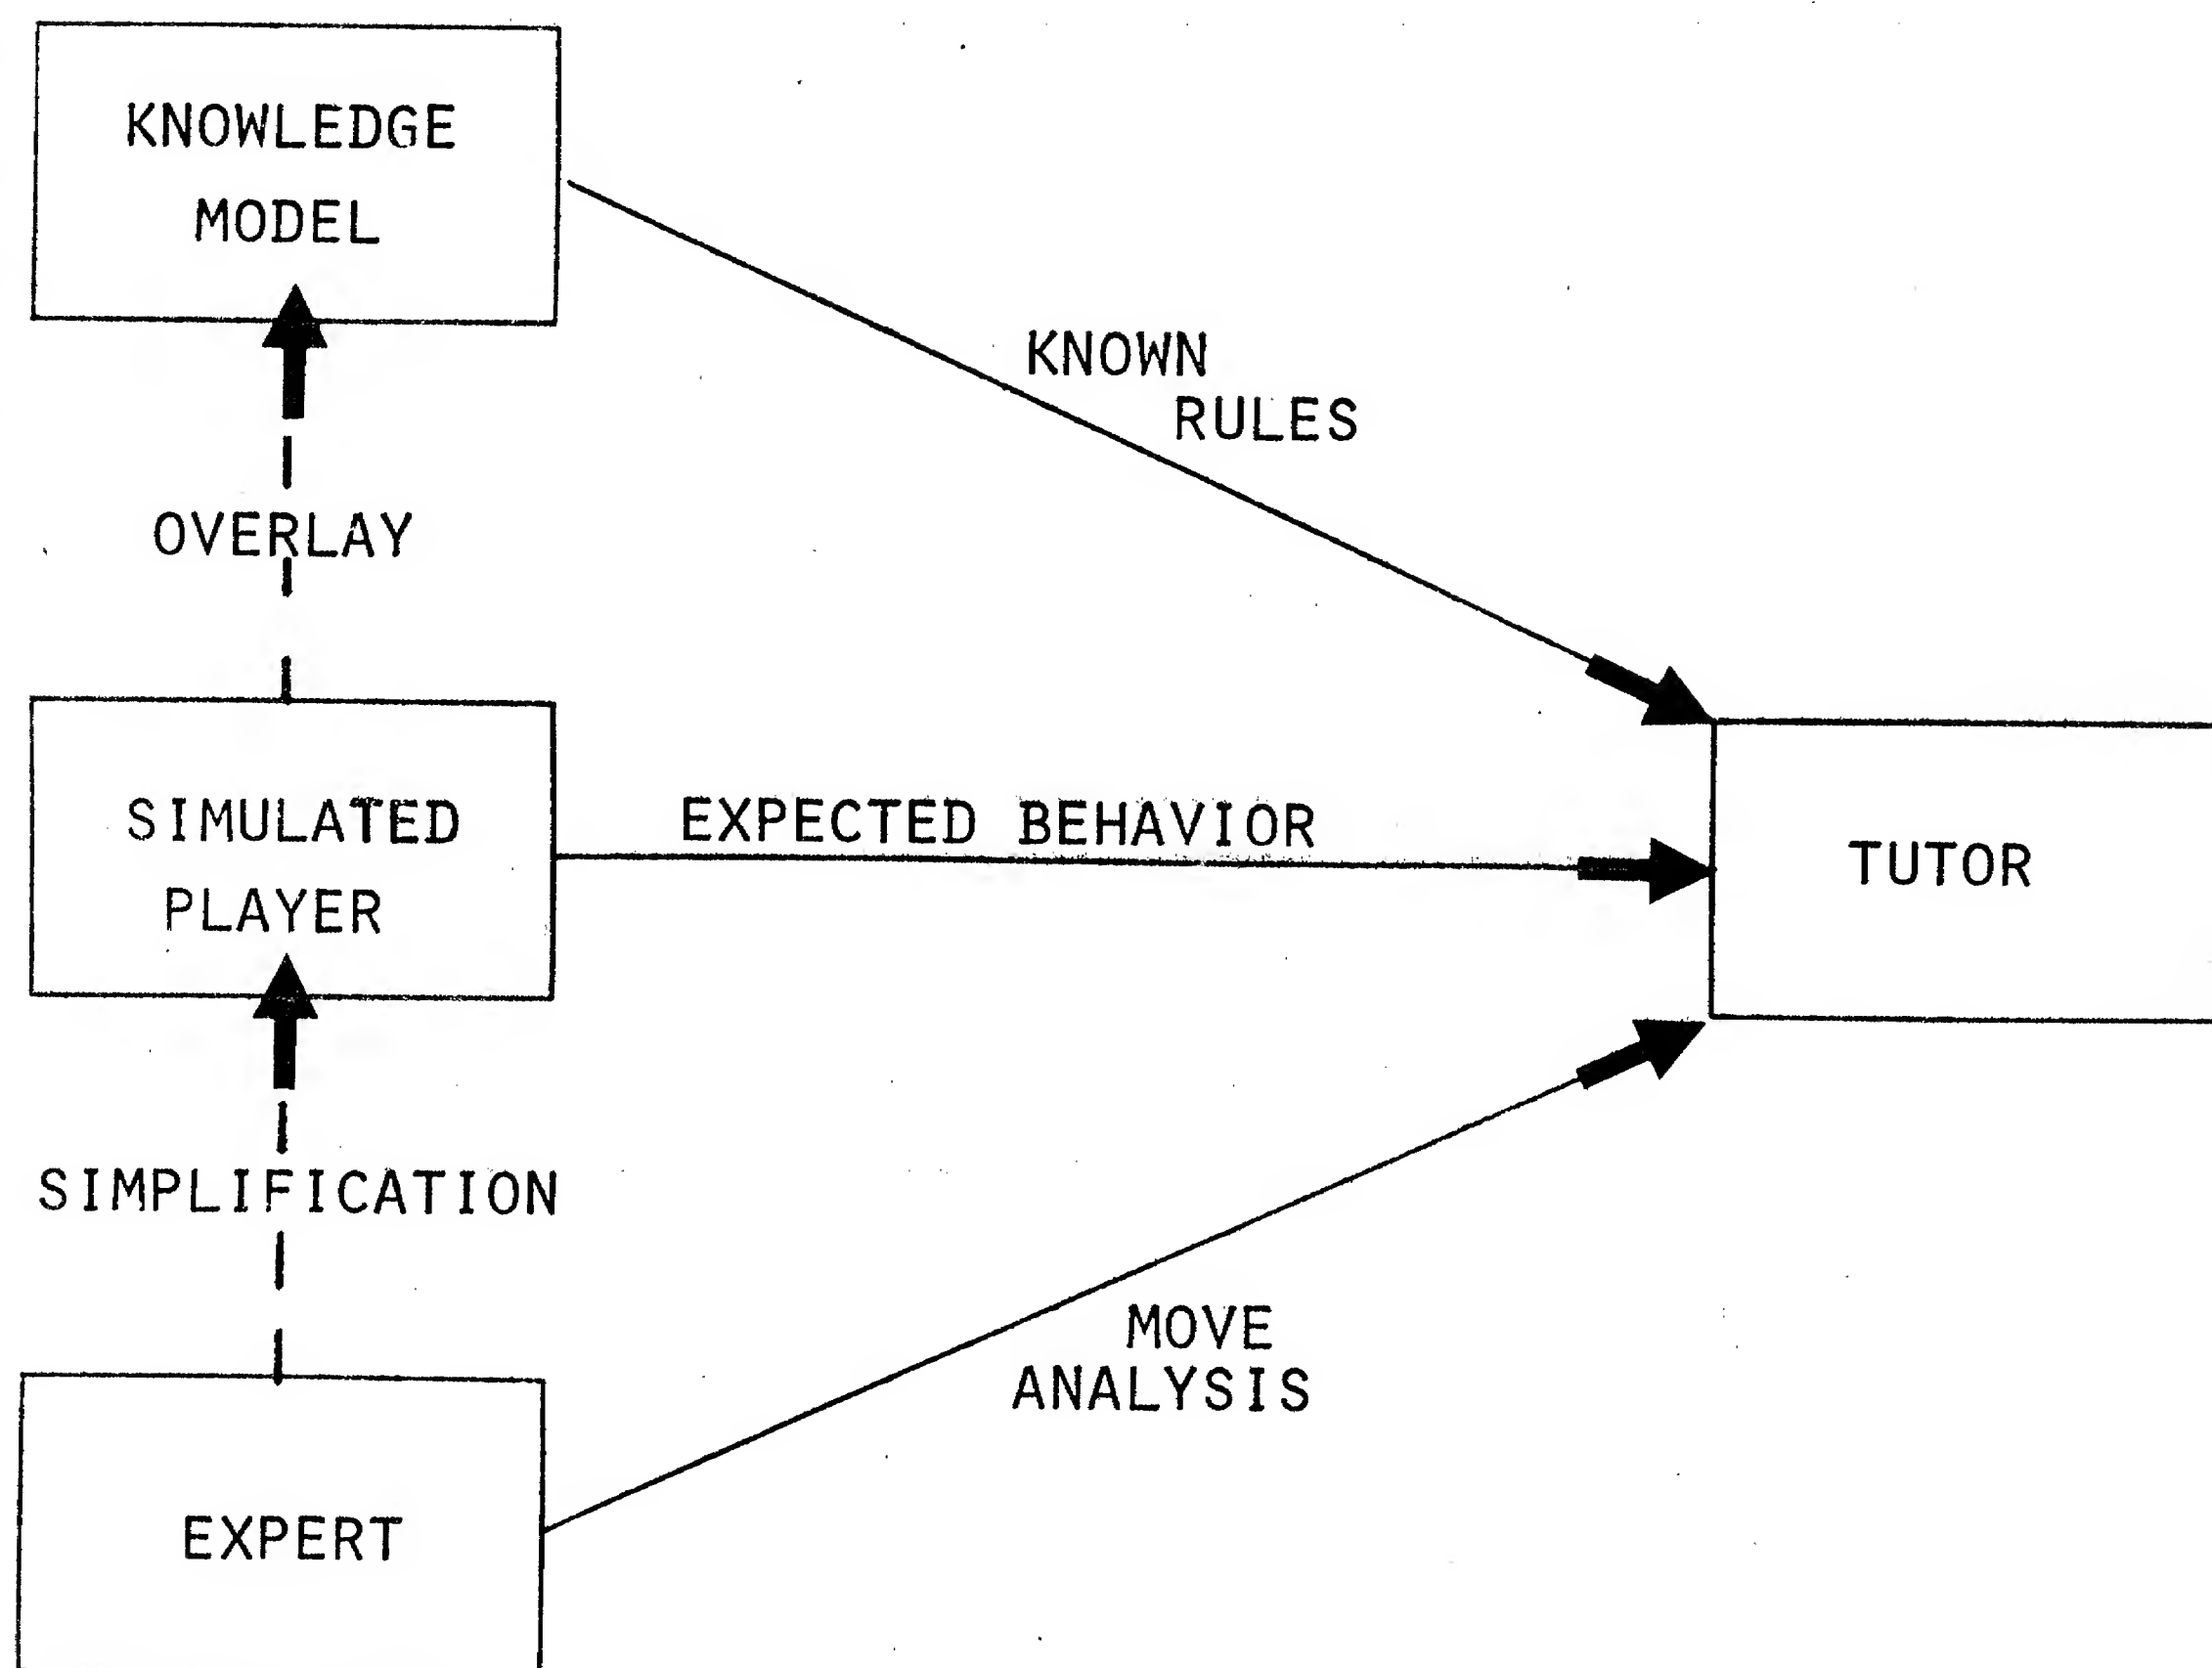


Figure 4

4. (dynamic dangers) Allowing the Wumpus to move, if the Player shoots at it and misses. (Recall that the game is won by shooting an arrow into the Wumpus' cave.) This requires the player to understand how old evidence degrades in this situation.
5. (competing goals) Playing the game under a time constraint results in the subgoal of visiting the nearest neighbor to obtain new information competing with the goal of always visiting the safest cave, no matter how far away.

A third use of simplification knowledge relates to a theory of player bugs. A player may possess a skill of the Expert in some approximate form that leads to bugs in certain situations. For example, the player may have formed an initial set of skills for Wumpus that never consider multiple dangers, but instead assume that only a single pit or bat is responsible for adjacent warnings. In the scenario, this would permit the incorrect deduction that cave 14 necessarily contained a pit. (See again figure 3.) Proper modelling requires the Coach to recognize that certain clauses are missing in the simplified logical rules possessed by the player. This can be done by checking whether a known simplification leads to behavior similar to the player's. In this case, the simplification discussed above involving approximations to rules obtained by dropping exceptions would produce a Simulated Player version of the Expert that mirrored this behavior. Thus, the simplification rules lead to a theory of "developmental bugs", namely those errors that arise due to rational simplifications of complex skills that arise as natural stages in the learning process. Sensitivity to this can allow the Coach to discuss the player's performance not as an outright error, but as a simplification inappropriate for more complex situations.

A fourth use of simplification rules is in summarizing a complex explanation for initial presentation to the player. This topic is pursued in the section on

tutoring.

To conclude, having a syllabus raises two issues. The first is whether a suitable sequence of simplified games might eliminate the need for the coach entirely. Our belief is that the choice of a simplified game facilitates learning by the player, but does not eliminate the need for the coach. The game cannot be made so simple that no difficulties for the player arise, else the game is no longer enjoyable; and, as long as such difficulties exist, the utility for tactful tutoring when they occur remains important. Furthermore, the simplified game makes modelling easier exactly because there are fewer skills potentially being employed by the player. Still we consider the utility of the Coach ultimately to be an experimental hypothesis that requires validation. One control on measuring the utility of the Coach will be a comparison of the rate at which skill is acquired with players who do not have access to coaching. (This is discussed in the plans for our experimental program in chapter 4.)

The second issue is whether the syllabus implies that all players acquire the skill in exactly the same fashion. This is not our approach. Rather we think of the syllabus as a general plan for the Coach that is adjusted for particular players on the basis of their current knowledge state. Indeed, knowledge and skills are not linearly ordered and we intend to structure the various constituent skills of the Wumpus expert into a partial order (on the basis of prerequisite constraints). It is the job of the tutoring module, in discussing issues with the player, to decide on what path through this network to take. This personalization is based on a consideration of the structure of the syllabus and of the individual knowledge state of the player. We discuss the Tutor below, but first an analysis of the pre-requisite modelling capability is required.



### 2.3 The Player Knowledge Model

Successful coaching requires a model of the player's current skill. For an intellectual game of the kind we are analyzing here, a model of the user's knowledge can be constructed as an overlay on the Expert, i.e. as a set of hypotheses regarding the relation of the player's knowledge to the Expert's. For each rule, the overlay provides probabilities that indicate the system's measure of confidence in the three alternatives: (a) that the player knows the rule, (b) that he knows the rule in some modified form (for example without knowledge of its exceptions), or (c) that he does not know the rule.

In some situations, it may be preferable to generate the Knowledge model as an overlay on a particular Simulated Player (as we showed in figure 4). For example, if the player is known to be a beginner, then typically he possesses knowledge of the logical and probability rules in a simplified form. So modelling can begin in terms of a Simulated Player with these characteristics. From this perspective, the Simulated Players can be viewed as average players at different levels of skill. This approach allows tutoring to take initial advantage of the coarse model provided by the Simulated Player, until a more detailed assessment of the player can be made.

The use of coarse models of average players was explored in our early implementations of a Wumpus coach and found to be a profitable first step toward personalization. A limited ability to make dynamic refinements of this initial coarse estimate using implicit evidence from the student's play is present in Carr's WUSOR program. However, a lexicon of Simulated Players representing the student at various points in the syllabus has not yet been studied, nor has overlay modelling been tested in this context.

### 2.4 The Psychologist

To generate and maintain a player Knowledge model, a modelling component in the Coach is required. We call this module the *Psychologist* (figure 5). This

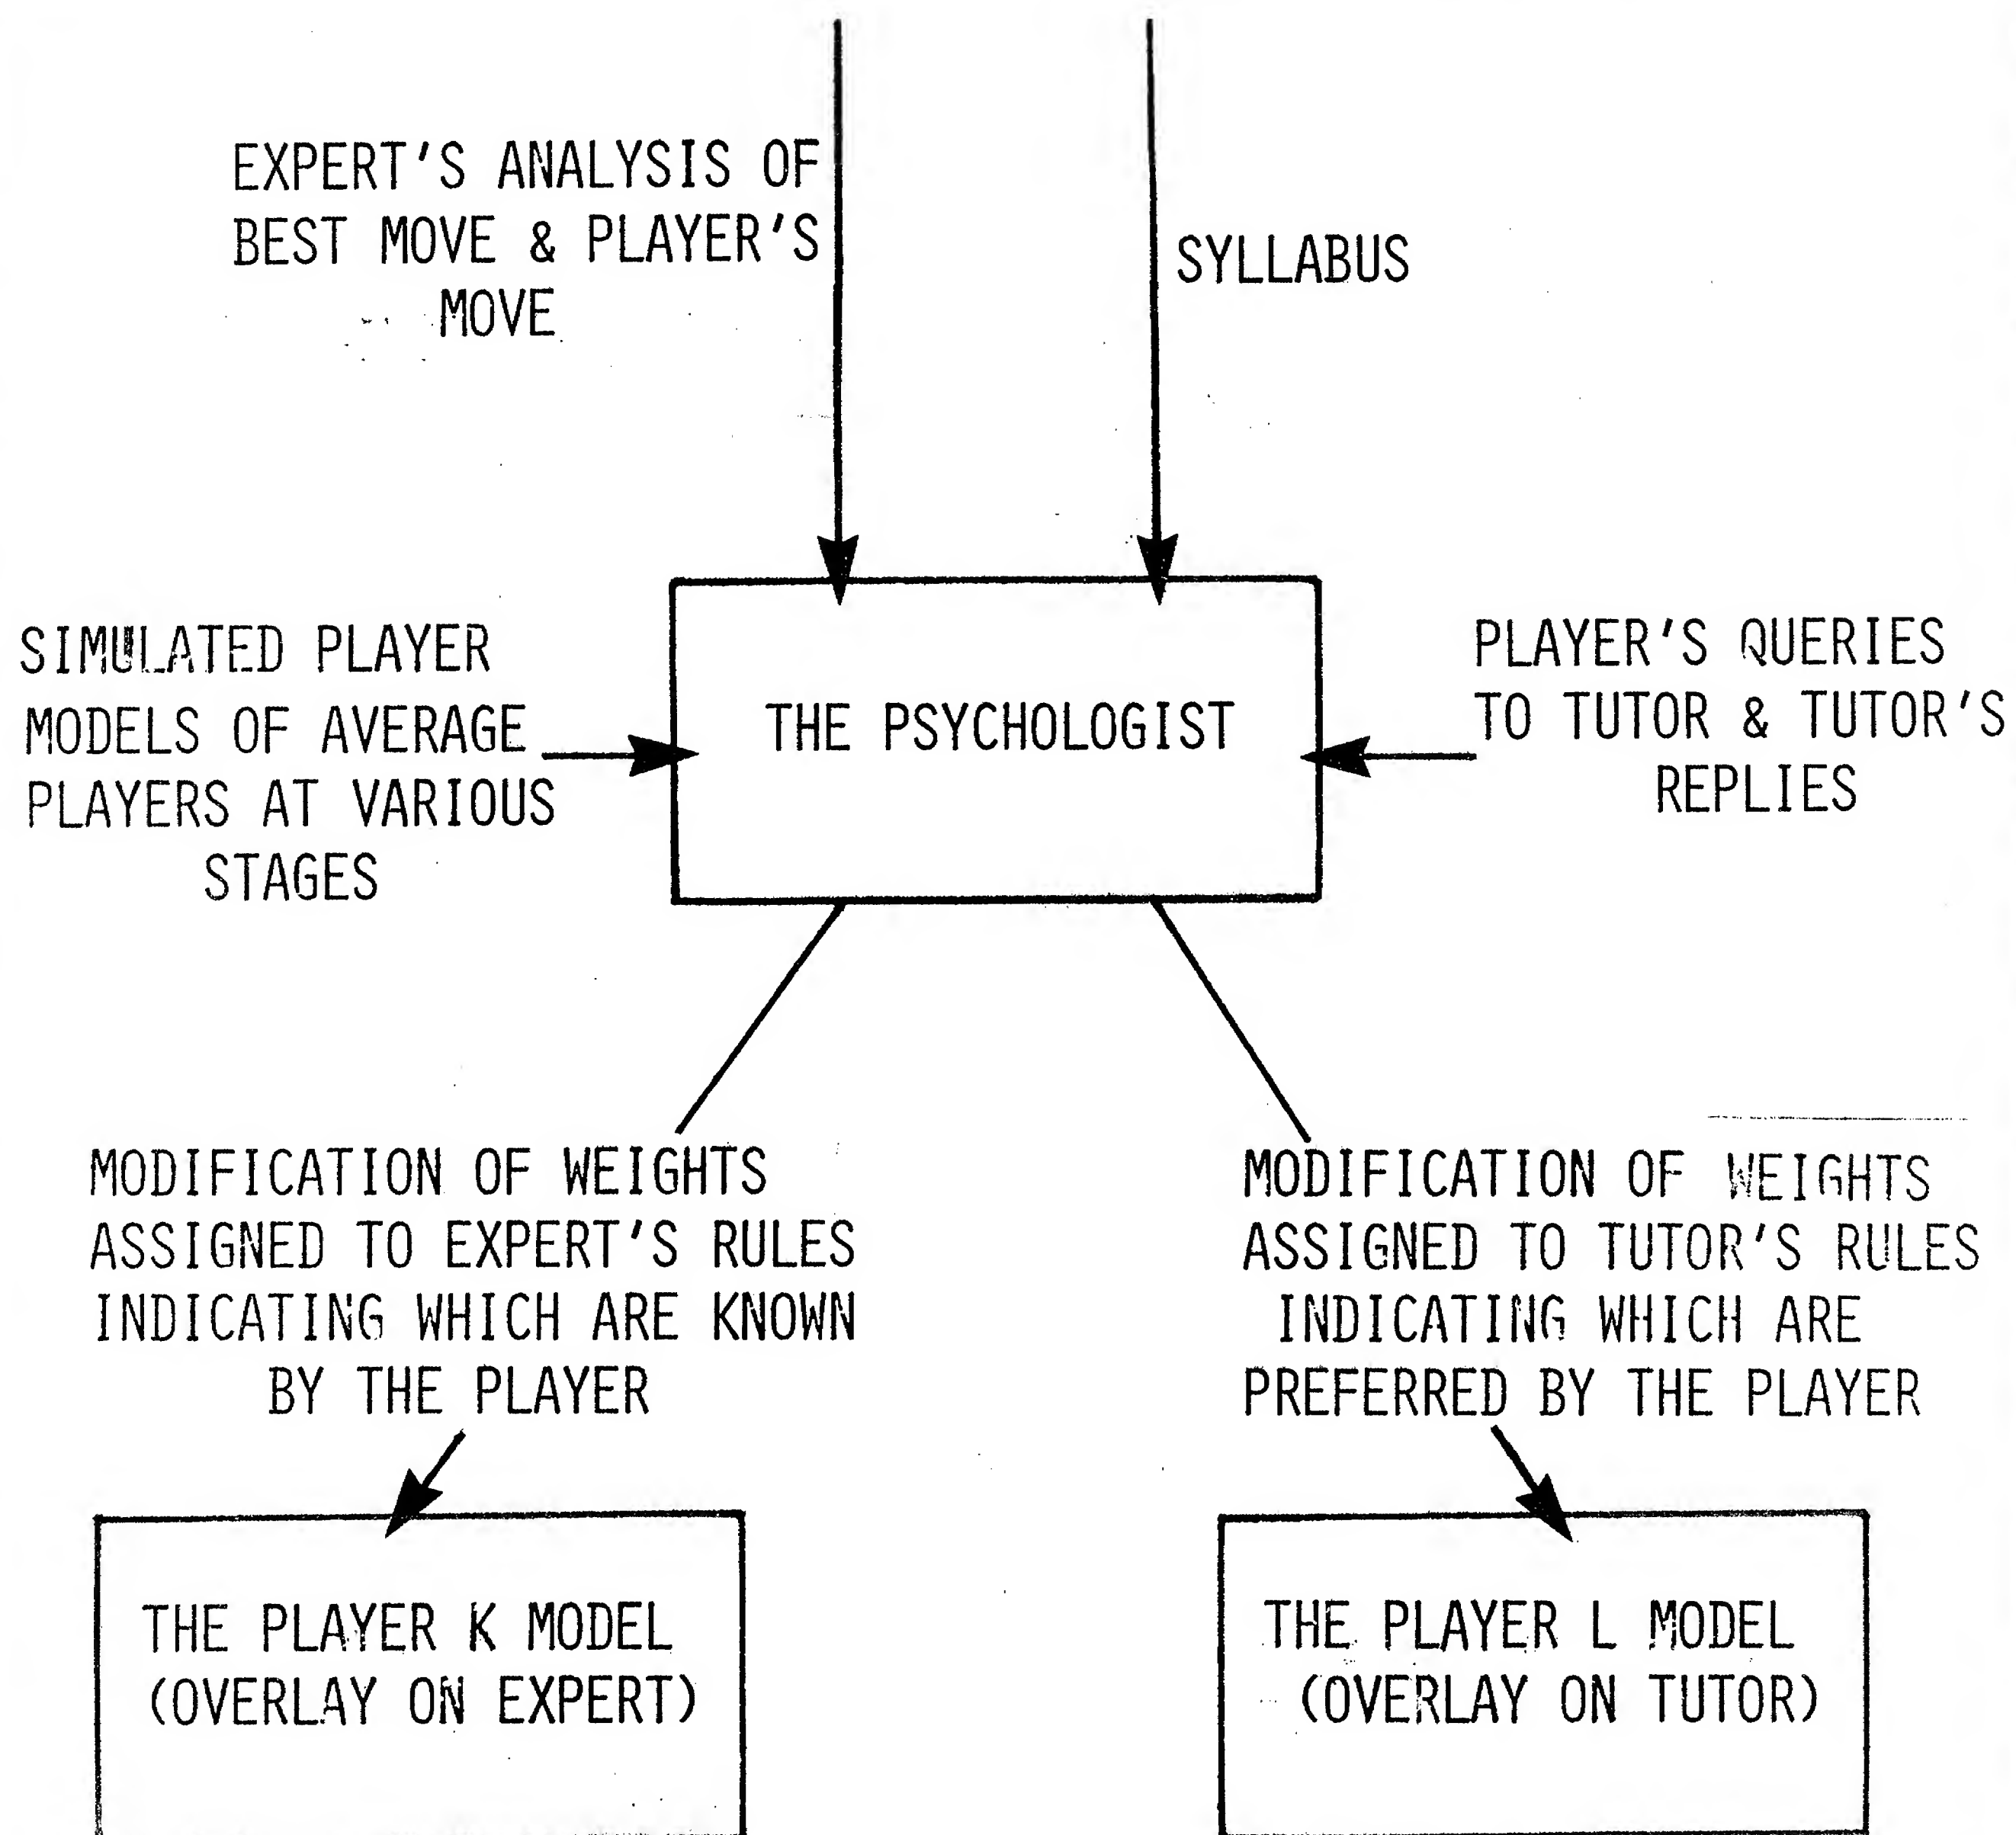


FIGURE 5 - THE PSYCHOLOGIST MODULE OF THE COACH

module is responsible for generating hypotheses regarding the player's knowledge on the basis of evidence arising from (a) the player's behavior in the game, (b) his explicit questions and directions to the Coach, and (c) the long term model of the player's position in the syllabus. We plan to implement and experiment with a modelling component using these criteria.

The fundamental kind of evidence arising from the player's behavior is the difference between his answers and those of the Expert. The Expert is able to explain what rules were involved in determining a given move. The Psychologist observes the difference in rules between those involved in the best move and those involved in the player's choice. The hypothesis is that the player is unaware of those rules which would be differentially involved in making the unrecognized better move. As an example, in the scenario given in the last chapter, the Psychologist would have considered the move to cave 14 evidence for increasing the estimate that the student is unfamiliar with the Double Evidence probability rule.

Such evidence can be misleading. The player may have known the rule, but his drawing of the warren was sloppy or incomplete, and hence he was unaware that the rule applied. Or perhaps the player has chosen to play quickly, and for that reason is not engaging in careful analysis. Hence, the Psychologist must deal in probabilities, not certainties. Furthermore, it will be important to allow the Psychologist to affect the version of the game being played, in order to create situations where it will be easier to decide if the player has a particular skill. For example, the Psychologist might direct the Coach to provide a display of the explored caves in order to insure that it is not a faulty sketch that is causing the difficulties.

Since we envision the Coach being used by a player over long periods of time, it will be important to include a *forgetfulness factor*, i.e. for the Psychologist to decrease the probability that the player knows a skill if it has not been

exercised over a long period. Or perhaps we may wish to have the Coach give the player "warm-up exercises" whose solution will yield evidence on how much the player has remembered. An example of such an exercise might be to present the player with figure 3 and ask him to list the three possible moves in order of increasing danger.

The narrow boundaries of the game makes it possible to approach these difficult questions. The design, testing and debugging of alternative modelling strategies will be an important dimension of this research.

### 2.5 The Tutor

Successful expertise and modelling will be of no avail unless the Coach is competent in delivering advice. While we cannot approach the kind of empathy that can exist between human player and coach, we can attempt to give the Coach some flexibility in its methods for discussing particular issues. This capability resides in the Tutor module of the Coach (figure 6). The Tutor uses evidence from the player knowledge and learning models as well as from the Expert to guide the tutorial interaction. How these sources of evidence are used by the Tutor and to what ends is the subject of this section.

There are three dimensions to generating an explanation which must be considered by a Tutor: these are (a) What to say; (b) When to say it; and (c) How to say it.

For the Coach, *what to say* is a function of the differential knowledge (called the *topic-set*) between the move the player made and the best move available. In the existing WUSOR program, the Expert generates a *rule-set* for each possible move. The rule-set consists of those rules involved in computing the probability that a given cave contains a danger.

This quality of being able to explain its decisions is so crucial to an Expert program being used as part of a tutoring



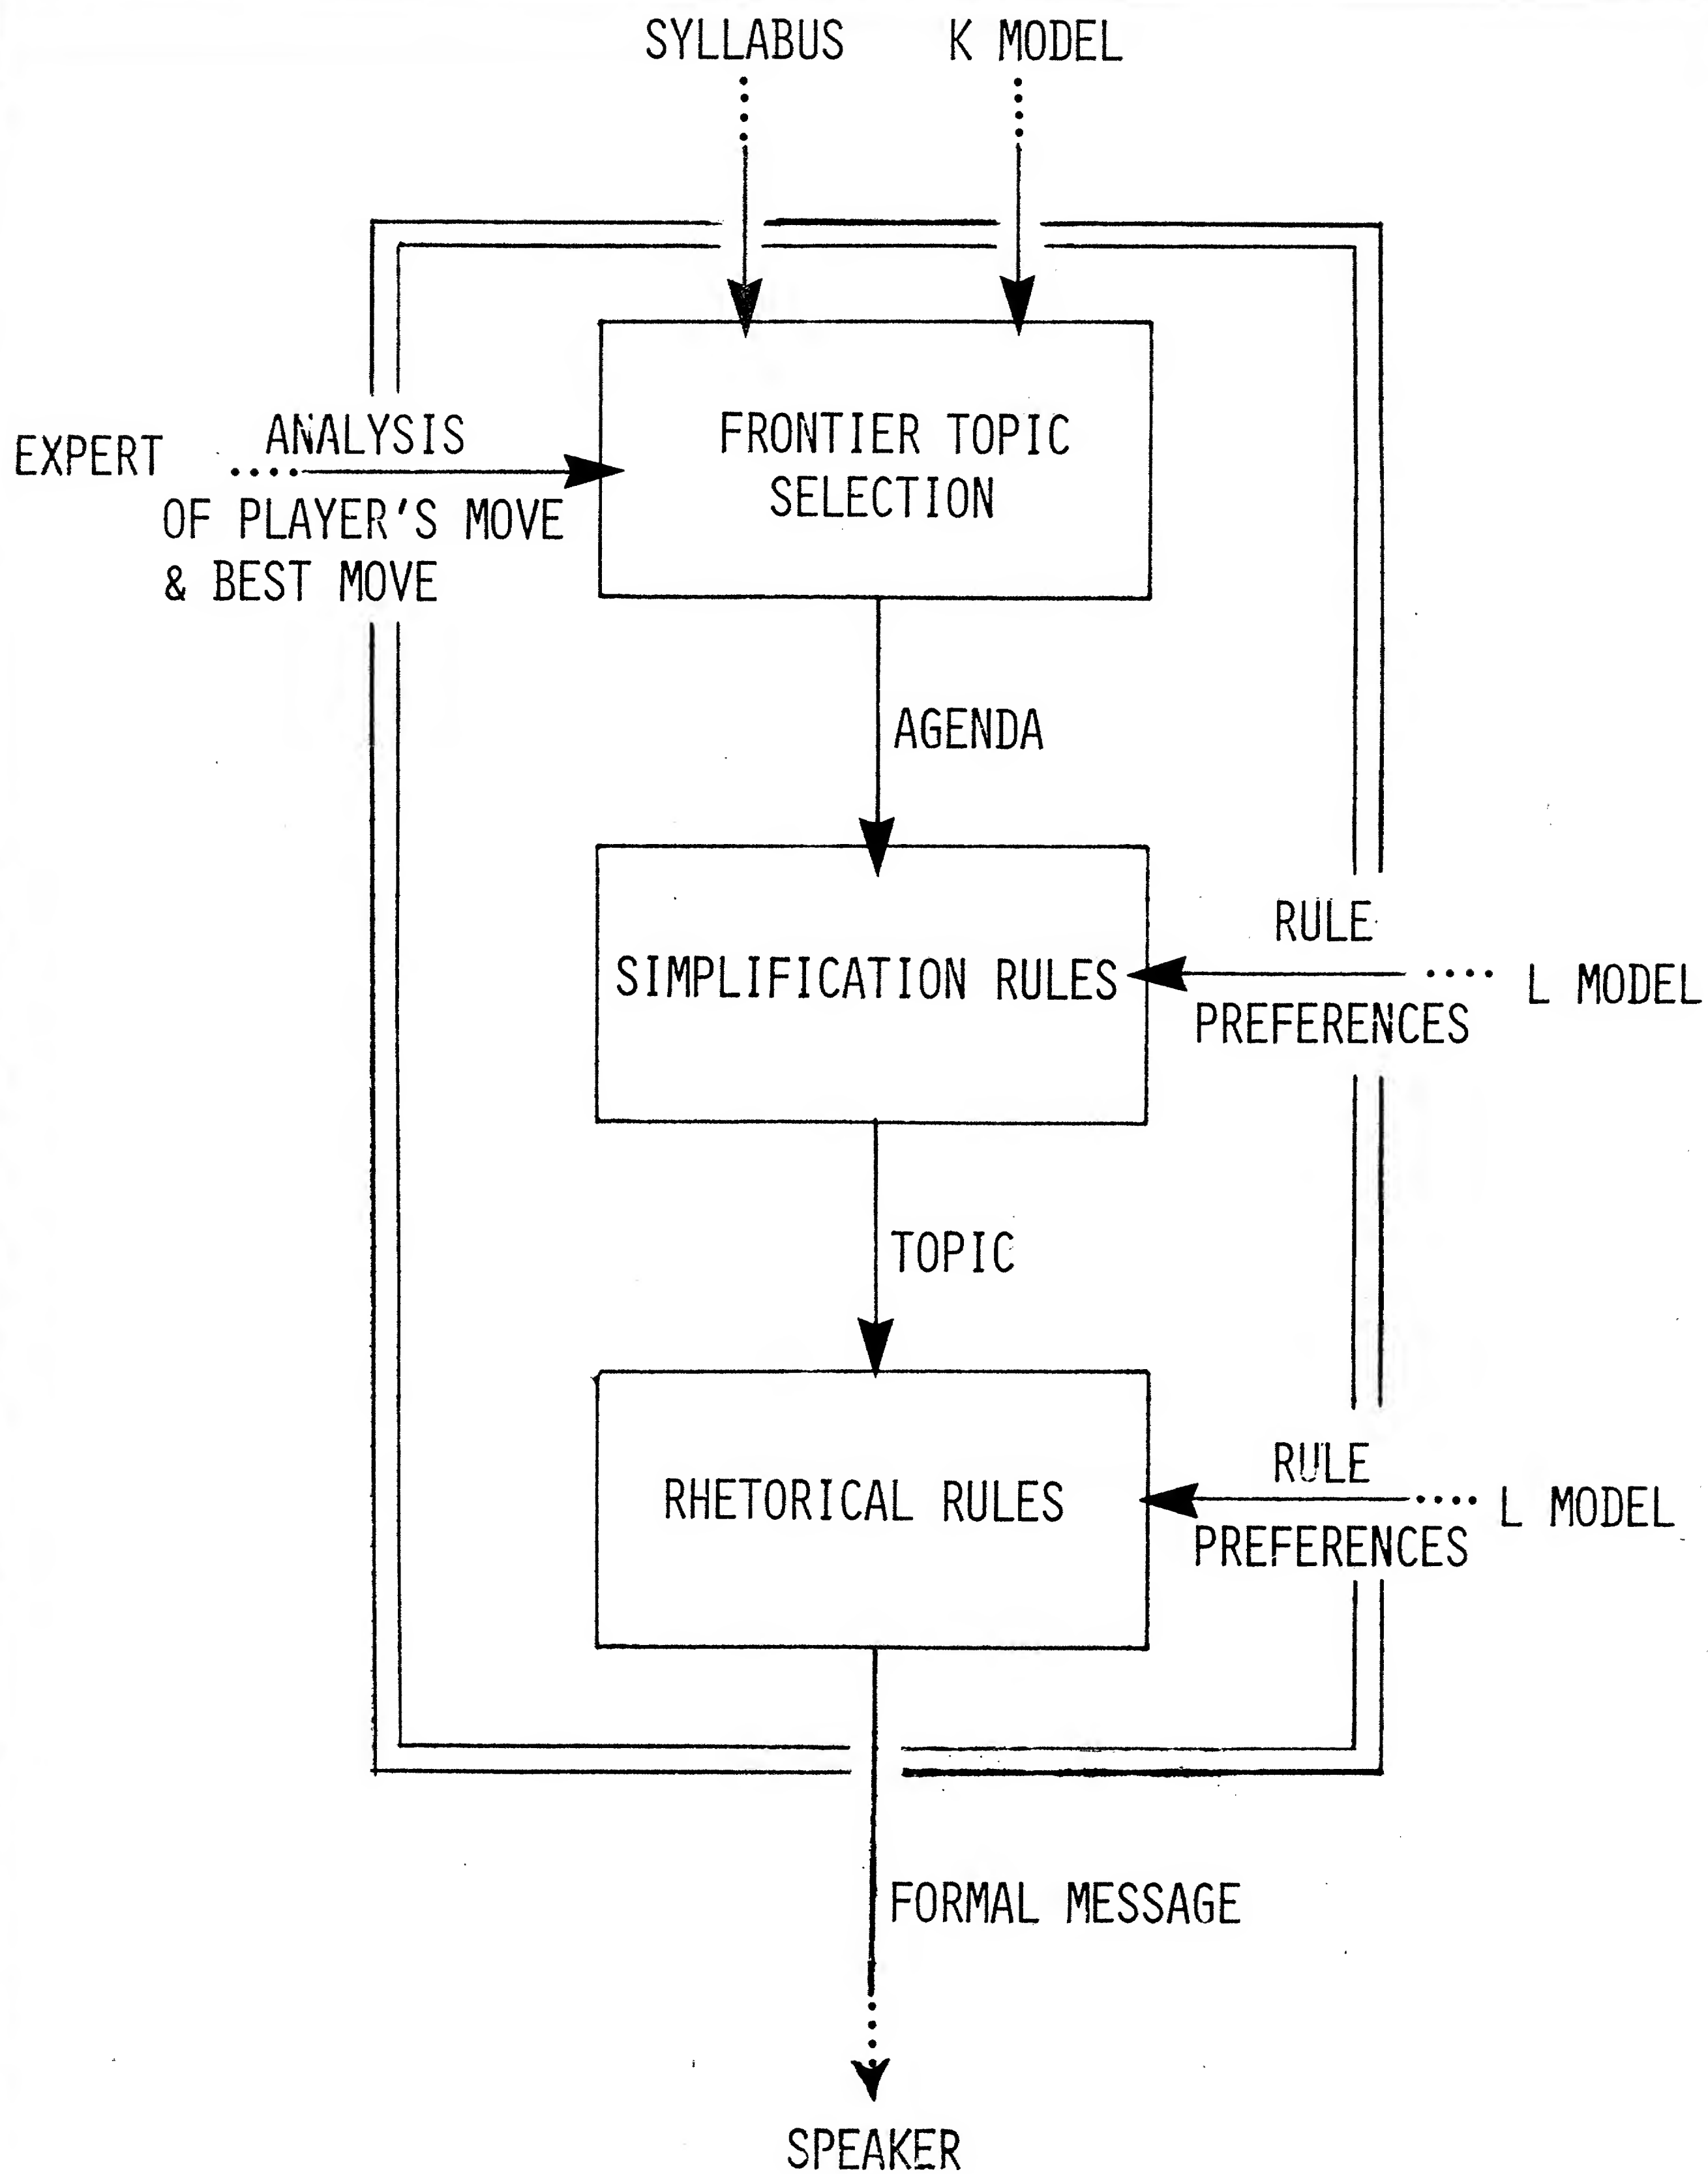


FIGURE 6 - THE TUTOR MODULE OF THE COACH

system that we call it the *principle of articulate expertise*.

We chose a rule-based approach to representing expertise because it has this "articulate" quality: the formal explanation of the probability of the danger of a given cave is a list of the rules needed to compute that probability. The *agenda* presented in the scenario was obtained essentially by noting the rules needed to determine the danger of the various possible moves.

The WUSOR Tutor then compares the rule-set of the best-move cave (i.e. the cave with the lowest probability of danger) with the rule-set of the cave chosen by the player. The difference between these rule-sets is the *topic-set*: the set of rules that are reasonable candidates to discuss with the player. The educational hypothesis here is: if the player knew the rules in the topic-set, then he would have been able to compute the correct probability of the best move, and hence would have chosen that cave. For the scenario, WUSOR would have generated a topic-set containing the double evidence rule along with several other topics.

WUSOR prunes this topic-set by eliminating those rules that its Knowledge model indicates are already known by the player. (It was on this basis that the logical rules were pruned by the Coach in the scenario). A complementary improvement is achieved by comparing the remaining rules in the topic-set with the Syllabus. A particular topic is chosen by selecting the rule of the topic-set nearest to the player's current position in the Syllabus. (In the scenario, this was illustrated by the Coach choosing not to discuss the choice of move quantitatively. Presumably the player was not far enough into the syllabus to merit this level of discussion.) Thus, as in human discourse, knowing "what to say" is improved by having a better understanding of your listener.

We call this approach to topic selection *frontier tutoring*. Figure 7 illustrates the origin of this phrase. The network of nodes represents the




FIGURE 7 - FRONTIER MODEL OF TOPIC SELECTION

syllabus, with concepts increasing in difficulty from left to right and linked by dependencies. The shaded region represents the subset of the syllabus known to the player, as indicated by the Knowledge model. The darkened nodes are the unpruned topic-set. The frontier consists of those nodes on the boundary of the shaded region. The preferred choice of topic are the topic nodes in the unexplored region of the syllabus closest to the frontier. COACH-1 will employ a similar strategy, augmented by the Simplification rules and Rhetoric rules outlined later in this section.

The coaching paradigm also suggests *when* to engage in tutoring. If the player might have made a better move, then he will often be interested in knowing this, in order to improve his play. Hence the Tutor engages in a discussion with the player about the underlying rules when their availability would make a difference to the player's decision. The Tutor's goal is to convey the underlying knowledge: the player's is to become an expert at a game he enjoys. In the Wumpus scenario presented earlier, this was illustrated by the coach using the move to the dangerous cave 14 as an opportunity to discuss the "double evidence" heuristic. This issue and example oriented tutoring fits within the paradigm introduced by Burton and Brown [1976] for an arithmetic game called HOW THE WEST WAS WON, originally developed for for the PLATO Elementary Mathematics Project by Bonnie Anderson.

However, the research proposed here extends Burton and Brown, (as well as the existing WUSOR program), by considering the procedural formulation of a broader set of explanatory techniques. For example, it is not enough to have an issue to discuss and an example to illustrate it. Care must be taken to control the length of the explanation. For this reason, we plan to make simplification rules available to the Tutor for summarizing explanations. Examples are:

*S1. Simplify "A because B" by A.*



S2. Simplify " $P(x) \Rightarrow P(a)$ " by  $P(a)$ .

S3. Simplify " $A \& (A \Rightarrow B) \Rightarrow B$ " by  $B$ .

These rules are suggested by work of Rumelhart [1975] on the summarization of stories. They were introduced earlier for creating Simulated Players by simplifying the Expert's rules.

To illustrate their application in this context, consider again the scenario. In our proposed design, the Tutor will generate a complete explanation internally, which we shall call the Agenda, and then apply the simplification rules to select the appropriate remark to be made to the player.

For example, in the scenario, part of the agenda for the cave 14 double evidence explanation is:

1. There is double evidence that cave 14 contains a pit.
2. There is single evidence that cave 0 contains a pit.
3. Single evidence is safer than double evidence.
4. Therefore, cave 0 is safer than cave 14.

We have given the Agenda in English, though it would have a formal internal representation in the Tutor. Also it is incomplete, since there is no discussion of the other dangers. Applying the simplification rules, the tutor's formal response to a move to cave 14 would be its conclusion: *Cave 0 is safer than Cave 14.* (In the scenario, the Speaker has applied ellipsis, generating the abbreviated reply: *Cave 0 is safer.*)

Of course, the player may desire more information. Hence, as we indicated in the scenario, it will be possible for the player to ask questions, with the tutor answering by supplying more of the Agenda. We call this approach to providing advice a *discourse theory of explanation* to emphasize that an explanation is not a lecture, but rather an interactive dialog between Coach and player. We return to this theme when we introduce linguistic discourse rules as part of the

Speaker's capabilities.

We have considered "what" and "when" to coach the player. There is still the dimension of "how" to formulate an explanation regarding a particular rule in the topic-set. Collins [1976] has formulated as procedural rules a variety of Socratic tutoring techniques. The research proposed here extends this taxonomy. We propose to develop procedural tutoring techniques that fall within all four of the following categories (although our lexicon of tutoring techniques for each category will certainly not be exhaustive):

1. *Logical Explanations*: the most common example would be citing a rule plus the current evidence and drawing the correct inference. The above agenda for explaining the danger of cave 14 was of this form.
2. *Hypothetical Explanations*: the use of a supposition in developing an argument would be allowed here. As an example, suppose the coach is trying to explain that cave 2, a neighbor of cave 1, is safe from bats because 1 contained no squeaks. This is, in essence, tutoring the logical rule, "If a cave does not contain a warning, then no neighbor contains a danger." A hypothetical explanation would be:

*Suppose cave 2 contained a bat. Then we would have heard a squeak in cave 1. But we did not. So cave 2 must be safe from bats.*


3. *Graphical Explanations*: this is really another dimension -- a logical or hypothetical argument can be given in English or via pictures showing states of the warren. Figure 8 illustrates a series of scenes on the TV display that parallel the preceeding hypothetical argument.
4. *Concrete Explanations*: this category would include explanations oriented

Figure 8

## GRAPHICAL EXPLANATION OF


"IF A CAVE DOES NOT CONTAIN A WARNING, THEN NO NEIGHBOR CONTAINS A DANGER."

1. SUPPOSE THE WARREN LOOKS LIKE THIS:




2. THEN YOU WOULD HEAR A SQUEAK.

3. BUT,



4. THEREFORE,



around specific examples rather than rules, leaving it to the player to make the necessary generalizations. The simplified initial coaching to the student in the case of the cave 14 example resulted in the concrete remark: "Cave 0 is safer than cave 14." Only upon optional subsequent questioning was the explanation raised from the concrete to the abstract, with general principles given.

In essence, these explanatory techniques are a procedural theory of *rhetoric*. The word has acquired a negative connotation of "undue use of exaggeration or display", but its classical meaning is "the ability to use language effectively". We have generalized rhetoric to apply to graphical as well as linguistic explanations, but otherwise we are engaged in the classical study of effective communication for the purpose of conveying an explanation. Without rhetorical skill, a tutor, whether human or machine, will not be effective. Our concern for multiple explanatory devices is one of the qualities that distinguishes this research from classical frame-oriented CAI.

Given a catalog of possible explanatory strategies, the Tutor must select a particular strategy for a player at a given stage of development. To accomplish this, the Coach employs a Player Learning Model.

### 2.6 The Player Learning Model

There are many dimensions to a player's learning behavior, but the slice that is pragmatically useful to the Coach is the player's preferences with regard to the available tutorial modes. The Player Learning Model is an *overlay* on the Tutor; specifically, four numerical weights are associated with the each tutorial strategy which estimate whether:

- (a) the player is known to prefer that strategy;
- (b) he is known to dislike it;
- (c) the strategy has been recently used;



(d) the strategy has been successful as judged by subsequent play.

We introduced the *Psychologist* module of the Coach earlier, describing its generation of the *K* Model. The Psychologist is also responsible for generating and maintaining the *L* Module. The kinds of evidence we plan to use are (a) whether a given tutorial technique on the average results in the player successfully acquiring the rule the Tutor is trying to convey, (b) the player's explicit reactions to a given tutorial technique (e.g. "I don't understand!") and (c) general knowledge about the relative success of different techniques for various kinds of tutoring situations (e.g. abstract explanations are probably not preferable for young players or those not mathematically trained).

The Tutor will adjust its explanation of a given rule on the basis of the Learning model, although this does not mean that the player's preferences are always followed. For example, the Learning model might reveal that the player has become too dependent, as evidenced by always demanding a complete explanation. In such cases, the Coach would choose to refrain from providing the complete Agenda, despite the student's desires.

Successful modelling of learning is dependent on successful modelling of the various knowledge states the player passes through. Hence, our initial research will focus on generating the overlay model of the player with respect to the Expert (i.e. the Knowledge Model) to be followed later by an investigation of modelling the player's interaction preferences as an overlay on the Tutor (i.e. the Learning Model).

## 2.7 The Listener

The communications module consists of two components: a Listener for translating questions by the player into a form comprehensible by the Coach and a Speaker for translating the selection of a formal topic by the Tutor into a comprehensible form for the Player (figure 9). In this section, we discuss the Listener, and in the next, the Speaker.

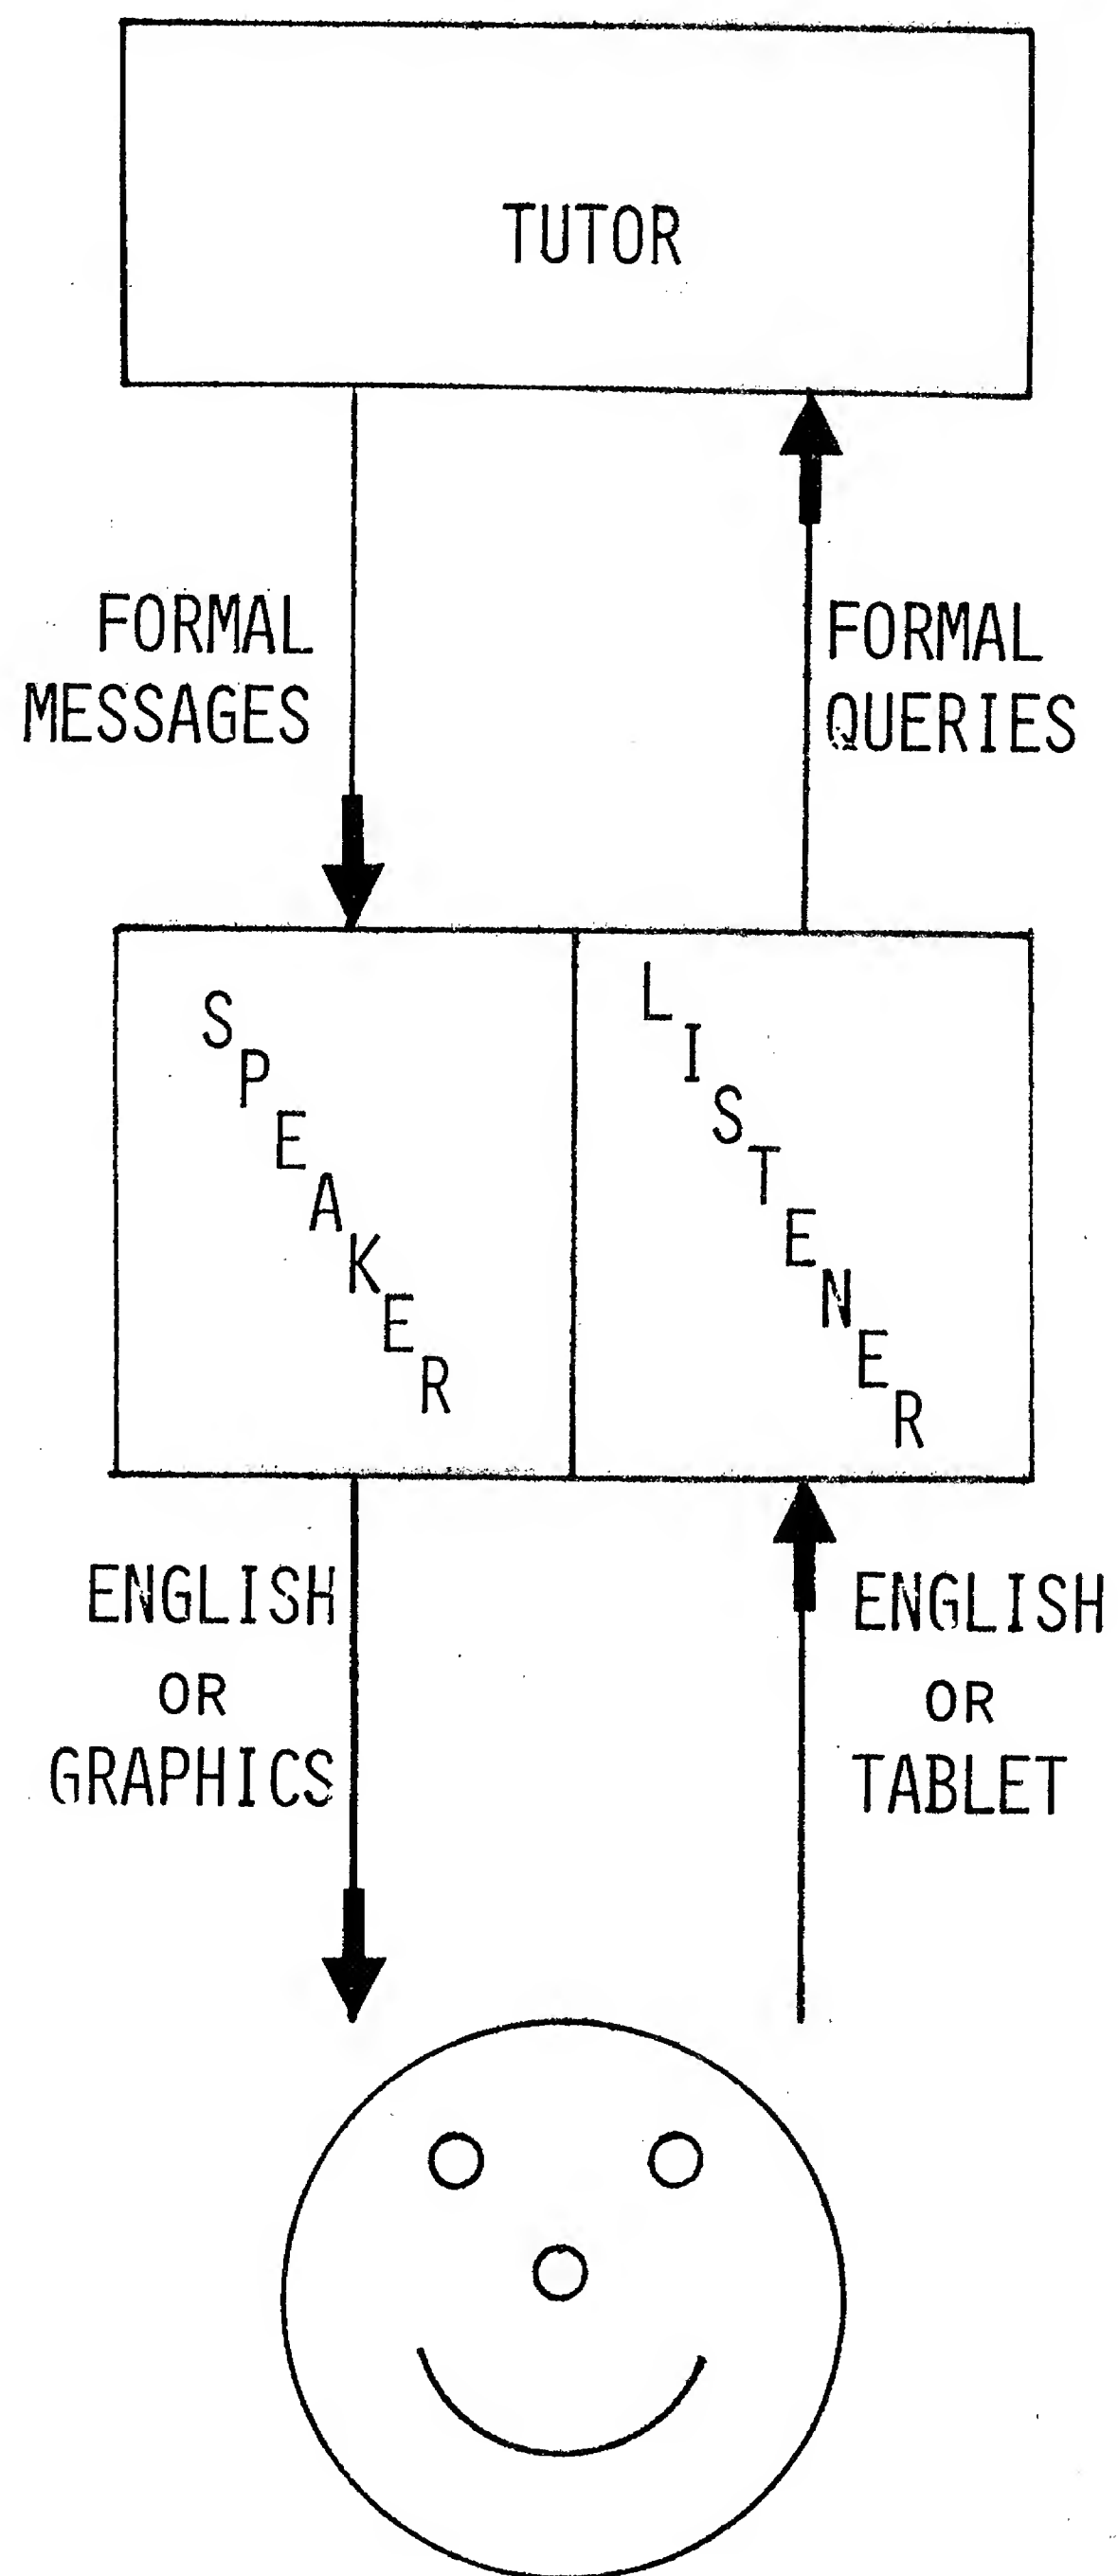


FIGURE 9 - THE COMMUNICATIONS MODULE OF THE COACH

Progress in AI offers the possibility of comfortable English comprehension by the Coach for restricted domains, as illustrated by the SHRDLU [Winograd 1972] and LUNAR [Woods 1972] programs. The use of semantic grammars in Sophie proved this for the electronics arena [Burton 1976]. We propose to apply this AI technology to the communications interface with the player, to the end of improving the system's sensitivity to the player's reasoning. Ideally, we would like to be able to ask the player for his explanation of why he made a particular move. How far along the spectrum of successful natural language comprehension we can reach is uncertain, but available AI technology will allow us to achieve qualitative improvement over previous computer assisted instructional environments.

There is another possibility for input from the player to the Coach -- the use of a tablet to support graphical input. We partly avoid the need for graphical input by having the Coach maintain a graphical display of the explored warren on the TV, but ideally the player should be allowed to maintain his own pictorial representation. If purely linguistic input proves too constrictive, we will add a tablet input module. Work by [Negroponte 1971, 1974; Herot, 1975; & Purcell 1976] indicates the possibility of achieving significant power in this dimension.

### 2.8 The Speaker

The Coach will also use English in its explanations. An ability to speak succinctly is critical if the interaction is not to be excessively tedious. We have already examined criteria by which the tutor can select the content of what it wishes to say. But after this has been done, there remains conversion of the conceptualization to English.

Our goal is to develop a "discourse oriented theory of explanation". Hence, we propose to embody in the Speaker various linguistic conventions for facilitating discourse. One such linguistic device which allows brevity is

*anaphora*, i.e. the use of pronouns to refer back to previously introduced objects. An example is:

*You visited cave 16 earlier. Cave 16 is a safe cave.*

reduced by anaphora to

*You visited cave 16 earlier. It is safe.*

Another linguistic device is *ellipsis*, i.e. not including a given phrase because the Coach can assume the player already has that knowledge. An example is

*You have been in two neighbors of cave 14 and you felt a draft in both neighbors of 14.*

reduced by ellipsis to

*You have been in two neighbors of cave 14 and felt a draft in both.*

Work on generation [Simmons 1973; McDonald 1976] suggests that it is a reasonable goal to include these capabilities in the Speaker in a modest way.

The Speaker module is also capable of generating graphic, as opposed to linguistic, explanations upon the Tutor's request. As the domain of discourse is a video game, this is essential. Thus, the term Speaker (and Listener) is being used in generalized form to apply to all communication channels between player and Coach.

## 2.9 Summary

To summarize our design for a Computer Coach, examine again figure 1. The Psychologist makes hypotheses regarding the player's state with respect to the Expert's knowledge and the Tutor's explanatory techniques. It does this using evidence from (a) the player's performance as analyzed by the Expert, (b) his explicit instructions to the Coach and (c) the expectations of the Tutor. These hypotheses form the Knowledge and Learning models. The Expert indicates when a



potential tutoring situation has arisen by informing the Tutor that a move is non-optimal. The Tutor uses the analysis of the expert and the content of the models to determine what topic is appropriate to discuss with the player and in what fashion. A natural language interface, the Speaker and Listener, serves to improve the communication between player and Coach. Simulated Players are simplifications of the Expert which more accurately model expected performance of the player.

The design of Computer Coaches is more than an exercise in computer programming: it addresses fundamental questions in education. For example, four major theoretical goals of this research are to develop and test:

1. an *articulate model of expertise* that supports modelling and tutoring by providing explanations of alternative decisions that can be made in a given task state. Our approach is to use a rule-based formulation of expertise, and use traces of the rules required to make a given decision as the explanation.
2. *overlay modelling* which describes a player in terms of the capabilities of the Expert or a simplification of the Expert (Simulated Player). Does the overlay model improve tutoring? Can predictions be made about the performance of the player? This goal is similar to that of Newell and Simon [1972] when they seek to construct production systems that model an individual. The new ingredient proposed here is the use of a coaching environment to obtain evidence and test success.
3. a *discourse oriented theory of explanation* in which a frontier model controls topic selection and a catalog of rhetorical techniques determines the form of the message. The Knowledge model, the Learning model, a net structured syllabus and an articulate expert all contribute to this function. A Speaker component converts a formal message into a user-comprehensible form, applying still further discourse rules to obtain conciseness. Progress in this

direction will be an important step towards a procedural theory of explanation.

4. *rules of simplification* that can be used to organize expertise into a syllabus, define restricted tasks for the player as exercises, and summarize complex explanations. The fact that this theory of simplification finds multiple applications reinforces our belief that the computational environment of the Coach touches deeply on central issues in the theory of pedagogy.

In exploring these goals, we are developing a general theory of coaching that applies not only to the player of a game, but, in general, to a novice engaged in any task wherein the skill can be modelled by rules.

There are, of course, limits to the psychological and educational questions discussed in this proposal. We have not considered, for example, such issues as player/coach empathy. This arises when we consider whether the Coach should be allowed to alter the location of dangers (consistent with the clues given to the player so far) to reinforce its advice. This ability might be useful when the Coach warns the player that a move is nonoptimal, but the player moves there anyway. If the danger was not actually in the high probability cave, should the Coach alter the game state so as to position the danger there? (For example, if the pit had not been in cave 14 in the scenario, and the player had moved there, should the Coach move the pit to that cave.) While this might reinforce the Coach's advice, it risks losing the player's enjoyment of the game. The Coach would probably be considered a "cheater" if the player knew of these rearrangements. But, should we allow the Coach to have a design that must be kept hidden from the student? We are not sure about the relative merits of this particular tutoring ability. (It can arise in an intermediate form if we allow the Tutor to advise the Game module on the positions of the dangers for the new game, but prevent any alterations once the game is begun. This would probabably

be considered fair by the player.) We plan to study the empathetic relation between Coach and player during the experimental phase of this research, and incorporate whatever insights are obtained into the design.

### 3. Relevant Research

Section 3.1 describes the computer coach as a natural evolutionary step in Computer Assisted Instruction, remedying some of its major defects of the past and drawing on some of its strengths. This step is made possible by:

1. The application of artificial intelligence theory and technology to the design for the coach (section 3.2).
2. The use of information processing psychology techniques to model problem solving in terms of rule sets (section 3.3).
3. The declining cost of computation (section 3.4).

#### 3.1 Computer Assisted Instruction

To see the place of the Computer Coach in CAI, it is useful to characterize CAI in terms of four periods (figure 10). The first, which we have labelled the *Primitive Period* predates computers and represents the original work with programmed learning texts. At the time this work was undertaken, both the technology and the cognitive theory were in a primitive state.

The use of computers initiated the next phase which we term the *Classical Period*. It occupied most of the sixties and even today remains the dominant paradigm outside the research environment. Programs developed within this era were typically organized as a decision tree of multiple choice questions, with the student's responses determining which path in the tree is taken. These CAI programs were the first explorations of the computer as an educational tool. They were in some cases able to provide interesting learning environments, for example the chemistry programs of PLATO [Bitzer and Johnson 1971], but were ultimately limited by a minimal understanding of the problem domain being taught, and minimal models of the teaching and learning processes. The paradigm of the



<u>PERIODS IN CAI</u>			
PRIMITIVE	CLASSICAL	ROMANTIC	MODERN
PROGRAMMED INSTRUCTION USING WORKBOOKS	PLATO  TICCIT	SCHOLAR SOPHIE SET THEORY WITH EXCHECK WUMPUS BIP	→
NO COMPUTER	NO DOMAIN- SPECIFIC EXPERTISE	USE OF AI FOR EXPERTISE	

Figure 10

classical period was to develop tutor languages to facilitate the design of scripts by teachers for their domains. Such an approach to CAI remains useful in certain contexts, but to achieve a new plateau of performance, a new design philosophy is necessary.

The *Romantic Period* represents the shift to a new paradigm in which the goal is to embed genuine domain expertise in the CAI program. Three benchmark efforts in this category, each concerned with a very different kind of expertise, are the Logic and Set Theory tutors constructed by Suppes et al.; the geography tutor of Carbonell and Collins; and the electronics troubleshooting tutor of Brown et al.

\* Suppes has been involved with CAI since its inception, and hence his work spans the classical and romantic periods. One of his long standing goals has been the development of a proof checker capable of understanding the validity of a student's proof. With the gradual evolution of AI techniques, he and his colleagues have been able to evolve successfully more powerful proof checkers [Suppes 1972, Goldberg and Suppes 1972, Smith et al. 1975]. Thus, in this case, the research represents an evolutionary rather than revolutionary transition from classical to expert-based CAI.

\* Carbonell designed Scholar around 1970 as a CAI system for geography that could answer as well as ask questions. The basic theoretical improvement was the use of a semantic net, a dominant AI representation, to represent domain knowledge. Since that time, Scholar has evolved as a result of the later work by Carbonell, Collins and others [Carbonell & Collins 1973, Collins et al. 1975]. The program has served as an impetus for improving the power of semantic nets, and hence has had important feedback into AI research.

\* Brown's SOPHIE system, a Sophisticated Instructional Environment for tutoring electronic troubleshooting, is impressive in terms of its level of domain expertise [Brown et al., 1975]. The program is capable of simulating the

internal behavior of a power supply, and hence can answer most student questions regarding the state of the device.

These programs, sometimes called Generative CAI [Bryan 1969], made possible a new level of performance. Such CAI tutors are not limited to comprehension of a highly restricted set of student responses; but, through an embedded domain Expert, are able to comprehend a much wider set of interactions. They were originally *romantic efforts* in that the AI technology of the late sixties necessary to implement Expert modules was itself in a relatively primitive state. But during the last six years a progressively more powerful set of AI tools have been developed and applied to embedding expertise in CAI programs. This evolution began with the original use of semantic nets in Scholar, followed by the use of semantic grammars and multiple representations of knowledge in Sophie [Brown & Burton, 1975]. Brown has referred to tutors of this kind as ICAI (Intelligent CAI).

Recently, a fourth phase in CAI research has begun, characterized by the inclusion of expertise in the Tutor regarding *the student's learning behavior and possible tutorial strategies*. We have chosen to call this new generation of instructional programs *AICAI Tutors*, to emphasize the use of AI techniques in the modelling and tutoring components as well as in the Expert module. Within this context, Collins [1976] has investigated computational models for Socratic tutoring strategies. Burton and Brown [1976] in a tutoring program called WEST have introduced issue oriented models of the student's knowledge, rather than simple records of right and wrong answers. Atkinson and others at the Institute for Mathematical Studies in the Social Sciences have examined the representation of domain expertise as a network in which tasks and their requisite skills are represented [Barr, Beard & Atkinson 1975]. In this research, the BIP system for tutoring the computer language BASIC, a model is maintained of the student's familiarity with various skills, and the next task posed to the student is done



on the basis of which skills are currently known.

The research we propose here falls within this AICAI paradigm, and represents an integrated investigation into tutoring and modelling, in which the modelling component is concerned not only with the student's knowledge but also with his learning preferences. Both the BIP and the WEST research consider only a model of the student's knowledge, not his preferences for interacting with the tutorial system. The catalog of tutorial techniques we shall introduce subsumes those considered by Collins.

Of course, this division of CAI into four periods is a simplification. Since its inception, there has been a desire to model the student accurately. The BIP research, which applies AI representation techniques to the syllabus and to modelling the student, does not incorporate a powerful domain expert (and hence is limited in the complexity of the problems it can allow the student to undertake and still comprehend the student's results).

The novelty of the research we propose here is that in a single system there will be significant domain expertise, a broad range of possible interaction strategies available to the tutor, and a modelling capability for both the student's knowledge and his preferred modes of interaction with the tutor. If this research is successful, the early promise of CAI as a personal, responsive learning environment for the student will be a met in a far deeper and more fruitful fashion than classical CAI of the sixties was able to achieve.

### 3.2 Cognitive Psychology

Over the last fifteen years, a new psychological discipline concerned with the formulation of computational models of cognition has evolved. Benchmark texts in the field are Lindsay and Norman's *Human Information Processing* [1972] and Newell and Simon's *Human Problem Solving* [1972]. Formalisms such as semantic nets and production systems have been used to construct procedural models of memory and of problem solving. This computational approach has proven valuable



in elucidating aspects of human psychology that were not adequately explained by the more static theories of the past.

We propose to apply these new concepts of cognitive psychology to the representation of knowledge and learning models of the student, and to the design and implementation of automated modelling components. The individual's problem solving behavior will be described in terms of his knowledge of a set of rules, where the rules include both basic facts and control knowledge (i.e. statements about when they are applicable).

We have already done preliminary work in this area: [Miller and Goldstein 1976a,b,c,d] describe the process of parsing programming protocols by means of a rule-based theory of planning and debugging.

### 3.3 Artificial Intelligence

Designing an Expert for a game is a traditional AI project. Samuel's [1963] checker playing program and Greenblatt's [1967] chess program have attained significant performance levels, both having beaten excellent human players. The notion of constructing a Coach to tutor a player in the skills of the Expert is a natural extension.

The goal of coaching, however, adds the design constraint that the Expert be so constructed such that its expertise is comprehensible to a player. For this reason, we have proposed a rule-based approach in which knowledge is represented as a modular set of rules. The traditional game playing programs were not usually of this form. However, there are a set of more recent programs that achieve significant performance, whose design does fall under this paradigm: e.g. Dendral [Buchanan 1969] and Meta-Dendral [Buchanan 1972] for Mass Spectroscopy and for learning Mass Spectroscopy, Mycin for medical diagnosis [Shortliffe 1976], and EL for circuit analysis [Stallman and Sussman 1976]. All of these programs perform at the level of human experts for their domain of expertise.

There are limitations to rule-based systems. Interactions between the rules, exceptions to the rules, and context dependency are all critical technical issues. Recent research by Davis [1976] and Goldstein [1976] addresses these limitations and develops an approach to providing meta-knowledge about the rules in machine understandable form.

A careful analysis of the virtues of rule-based systems (comprehensibility, modularity, extensibility), their limitations and corresponding extensions appropriate to handling those limitations goes beyond the scope of this proposal. But the impressive performance of existing AI systems suggests that it will be fruitful to apply this technology to the design of Computer Coaches.

### 3.4 Computer Science

Ten years ago, a large computer installation such as one based on a Digital Equipment Corporations PDP-10 with KA processor would have cost \$1,000,000. Today, a LISP machine now under design in our laboratory will provide three times faster computation for one twentieth the cost. The recent NSF conference on "a ten year forecast for the impact of computers on education" was in unanimous agreement that the cost of computation will continue to decrease dramatically over the next decade [HUMRO 1976].

Our proposal to develop Computer Coaches will demand significant computational resources. Such resources are available now in the research environment, and economically feasible for schools within the next ten years. The theory of computer-based coaches (and its implementation as software), not the hardware technology, is the critical limitation.

#### 4. A Two Year Research Program

We propose a three phase research program for evaluating the merits of the Computer Coach paradigm. Since phases II and III are contingent on the success of phase I, support is requested only for the first phase as a two year project. The major portion of this chapter describes the experimental program for phase I. A summary of phases II and III is given at the end of the chapter.

##### 4.1 Phase I: A Computer Coach for Wumpus

Since an articulate Expert for Wumpus already exists, we can directly immediately focus our attention on the critical learning and teaching issues involved in designing improved modules for the Psychologist, Tutor, and Speaker. We estimate that the design and implementation of a complete COACH-1 will take approximately 10 to 12 months. The details of this design have been specified in detail in chapter 2.

Across the second year, we will carry out an extensive testing program to evaluate the success of the rule-based modelling and tutoring capabilities of COACH-1. These experiments and the proposed subject populations are described below. The experiments will serve to rigorously test our hypothesis that a Computer Coach can tutor mathematical and scientific knowledge in personalized and responsive ways, and that the skills acquired by the players are transferable.

We will undertake two basic categories of experiments. The first category ("global experiments") is concerned with the overall success of the Coach as a tutor of transferable intellectual skills. The second category ("local experiments") is oriented towards analyzing alternative designs for various modules of the Coach. The first category provides the critical measure of success for the overall research project; the second category suggests possible refinements to the design.



#### 4.2 Global Experiments

We shall carry out two major global experiments, one at the 12 month and one at the 18 month point in this project. Each global experiment will address two questions:

1. Does the coach facilitate the acquisition by the student of the intellectual skills needed to play Wumpus?
2. Are these skills transferred to other tasks?

The skills in question are the basic logical and probabilistic inference rules needed to make reasonable hypotheses about facts given uncertain data.

First we consider question 1: measuring the success of the Coach in tutoring Wumpus specifically. We will examine the performance of three populations of players on a common sequence of 25 Wumpus games. The first population will be uncoached, the second coached by the computer and the third coached by a human teacher. The sequence will involve approximately 10 sessions, each on the order of 45 minutes and extend over several weeks. Transcripts of each player's performance will be obtained, with statistics computed of his or her success. (How often does the player win? How efficient are his moves?)

Our hypothesis is that the computer coached population will acquire skill at the game faster than their counterparts who receive no tutoring and equal to the rate of those players given human tutoring.

We will also correlate the performance of members of the three populations with their skill in traditional mathematics. This skill will be measured by the standard achievement tests taken by the players. This correlation will indicate whether mathematical games such as Wumpus and their respective skills are accessible to student/players who have not been high achievers previously in this domain.



Question 2 examines whether those players in the three populations who have mastered the skills of Wumpus are able to transfer these skills to other domains. We will obtain evidence for transference by exposing the same three populations of students to (a) games involving similar skills and (b) problem exercises.

An example of a different game is "Clue". In this popular board game, the task is to identify the criminal from a population of suspects. As the detective, you are given various clues. The same sorts of logical and probabilistic inference rules (as are required in Wumpus) are applied to estimating who had the best motive, opportunity and means. Isomorphs of this game can be constructed for war game situations where you are given evidence about the location of your opponent's ships, and your task is to estimate their actual location (this is a traditional game called BATTLESHIP); for prospecting for gold; and many other situations. The ease of creating problems requiring the same skills is a consequence of the importance of the abilities required by Wumpus.

Problems that exercise the mathematical skills (logical and probabilistic reasoning) required by Wumpus will be chosen. We shall examine standard IQ and achievement tests as well as mathematical and probability texts to select these problems. We will construct a set to be administered as both a pre and a post test to the three populations.

We shall run several series of the "global experiment" on populations of different backgrounds. Our hypothesis is that the Coach will equal a human teacher in facilitating the acquisition of transferable skills in logical and probabilistic inference, as measured by performance on the problem set.

We do not claim that the *general* style of reasoning employed by a player will be affected by the Coach, given the limited time during which players will be exposed to the Coach. Our hypothesis is only that the *particular* skills tutored by the Coach are transferred. The possibility exists that the Coach could serve

as a general model for the student of a certain kind of mathematical reasoning, but this question must await successful completion of the first phase of research outlined here.

#### 4.3 Subjects

Education majors in an undergraduate college will provide the initial subject population. We will select a group with a range of achievement in mathematics. Given that these subjects will be future teachers, their opinions and advice regarding the behavior of the coach will provide important feedback for us towards improving its design. Our experience with the early WUSOR coaches indicates that the game is complex enough to be of interest to such students. Following an initial round of experiments with this population, we will work with students of the same age from different backgrounds (for example, chosen from a two year technical college) and then with secondary school students. The result will be evidence for the relative success of alternative teaching and modelling techniques for student populations of different levels of skill, age and background.

Each run of the global experiment will be done with 30 students, 10 in each population. Given several runs of the global experiment by the end of this two year program, sufficient evidence should be available to indicate clearly whether a more extensive evaluation of computer coaching is justified.

To reach a satisfactory design for the Coach, we also plan "local" experiments that analyze the performance of individual modules of the system. Three kinds of local experiments are defined -- AI, psychological and pedagogical. These experiments are described in the next three sections.

#### 4.4 Local AI Experiments

These experiments test whether the modules of the Coach perform successfully on certain highly controlled exercises that are necessary (but not sufficient)

conditions for their success in the real educational environment. Examples are: (a) at what level of skill does the Expert module perform the task compared to human experts, (b) if a Simulated Player is created by modifying the Expert, will the Tutor successfully diagnose the modification. To illustrate this second kind of AI experiment, consider a Simulated Player for Wumpus created by writing a program that has access to the logical and strategic rules, but can make no probabilistic inferences. The question then is whether the COACH-1 modeller can successfully diagnose this weakness. These AI experiments involve careful analyses of the capabilities of the modules involved. They yield hard facts about performance in certain settings but no psychological data. The next class of experiments addresses psychological issues.

#### 4.5 Local Psychological Experiments

These experiments investigate the relative success of alternative modelling components in estimating a given player's state of knowledge. We will use the same testing program as that outlined for the global experiment (Pre/Post testing on the game and its isomorphs), except that interviews with the students wherein they describe their rationale for various moves will be obtained. The evidence for the Knowledge and Learning models will be the extent to which the student's description of himself and his performance on isomorphs matches the hypotheses of these models.

There is also another form of validation. The Simulated Player whose rule set most closely matches rules attributed to the student by the Knowledge model can be run to predict likely moves by the student, in a given state of the game. The accuracy of these predictions with the student's actual move is evidence that the Knowledge model is reasonable.

Finally, we will supplement interviewing of individual players by recording the verbal interactions of several players working as a team. Brown, Rubinstein, and Burton [1976] in recent tests of the Sophie system have used a team



environment quite successfully to obtain insight into the analyses being made by the students. The players naturally explain their reasoning to each other, decreasing dependency on interviews by the experimenter.

#### 4.6 Local Pedagogical Experiments

Computer Coaches embody a theory of the syllabus and of alternative tutoring strategies for conveying that syllabus. We plan experiments that test alternative syllabi and alternative tutoring strategies. For example, with respect to the issue of when to interrupt the player, the Coach's behavior can range from one extreme of always discussing better moves if they exist to the other extreme of only interacting with the player upon explicit request. Our goal is to obtain evidence for mechanisms by which the Coach's tutoring component can dynamically alter its choice of interaction mode and tutoring topic. Thus we are not interested in the statistically best syllabus or teaching style, but rather what improvement over the statistical choice can be made by a Computer Coach that personalizes the form and content of the interactions on the basis of evidence available in the knowledge and learning models. These experiments will again involve relatively small populations, but involve extensive data-gathering of protocols and interviews for these populations.

We also intend a rather novel experiment -- as the Coach is able to explain itself, we propose to allow skilled human teachers to pretend to be student players and then to interrogate the Coach on its rationale for various tutoring remarks. If the Coach is successful, it should be able to reply to the teacher's queries in an acceptable fashion.

Finally, we will perform careful attitudinal studies of the subjects regarding their opinion of the Coach. Do they find it helpful? Does its intervention increase or decrease their enjoyment of the game? Do they find it too cryptic, too verbose, or appropriately concise? These studies will be made by means of questionnaires and interviews.



If these experiments -- both local and global -- yield evidence that the WUSOR-II Coach provides successful modelling and tutoring capabilities, then we believe the following two phases of work would be justified.

#### 4.7 Phase II: Experiments in Other Domains

In order to avoid errors arising from the examination of this paradigm for only one domain, we would undertake to implement Computer Coaches for other games. This effort would evolve through the same steps as the experiments outlined above. The result of this parallel effort would be a more solid set of evidence on which to build the basic outlines of our theory of Computer Coaches and on which to evaluate their success.

Our criteria for choosing one or more parallel domains are (a) that the game exercise basic intellectual skills, (b) that the design of an Expert be feasible, and (c) that the game be enjoyable and motivating. STEVEDOR is a possible candidate. Recall that in this game, the player is asked to load a cargo, given various sets of simple machines. The machines have costs associated with them: the task is to find the cheapest combination of simple machines adequate to move the weight to the desired location. Successful play involves in a natural way knowledge of elementary physics with obvious opportunities to tutor this subject matter. The game is simple enough to build an rule-based Expert. Our hypothesis is that this would be the only module to be effected by the change in domain.

Layman Allen has demonstrated the possibility of creating interesting intellectual games such as WFF'N'PROOF and THEORIES'N'QUERIES in the far more restrictive setting of a non-computer technology. With the availability of the dynamic capabilities of computers and video, the possibilities are unlimited. In phase II, we would select several games that serve to exercise important intellectual skills but differ in interesting ways from Wumpus, after a preliminary examination of the various candidates.

#### 4.8 Phase III: A General Theory of Computer Coaching

The experience of designing, implementing and evaluating Coaches for several domains should provide sufficient experience to develop a general theory of tutoring and of modelling. In the third phase of this research, our goal would be to articulate this theory and to redesign the Coaches constructed in the preceding phases to take account of these insights. The theory would contain criteria for formulating domain expertise as rule systems, for creating simplifications as tutoring goals, for modelling knowledge of these rule systems and for tutoring them. The same class of experiments would be undertaken. Positive results would yield clear evidence that a new kind of educational environment can be provided -- consisting of computer games and coaches -- that nurtures the development of transferable mathematical and scientific skills. Furthermore, since the theory is domain-independent, it will apply to traditional computer-based learning environments as well. Finally, the general theory, being a precise formulation of modelling and tutoring skills, is a candidate for a more rigorous theory of human teaching.

## 5. Resources

### 5.1 The MIT Artificial Intelligence Laboratory

The MIT AI Lab is a leader in artificial intelligence research. Intelligent Coaching Programs will not succeed without a heavy infusion of AI expertise in the Expert, the Tutor and the Psychologist modules of the Coach. The MIT AI Lab has faculty and graduate students who can supply that expertise.

Specific resources of the laboratory relevant to this research are: (a) expert programs for various domains (e.g. mathematical theorem proving, calculus, electronics, decision making); (b) natural language systems for both generation and comprehension; (c) advanced problem solving languages (e.g. Lisp, Planner, Conniver, Scheme); and (d) a powerful timesharing system with editing, and debugging capabilities that facilitates the rapid development of prototype programs.

### 5.2 The MIT Division for Study and Research in Education

MIT is concerned with the application of technology to education and, as a demonstration of that interest, has created the Division for Study and Research in Education. Psychologists and professional educators are on the staff of the division and will supply useful criticism of the experiments planned here.

### 5.3 Technology Transfer and Lisp Machines

The MIT AI Laboratory has developed a stand-alone computer that can provide at a reasonable cost the kind of computer power needed for these experiments and prototypes [Greenblatt 1974]. It makes the research feasible now, and serves as a vehicle for the practical dissemination of such coaching programs in the schools of the 1980s.



#### 5.4 The Logo Project

The Logo Project, a research group in the MIT AI Lab, is concerned with the development of improved educational environments based on the use of advanced computer technology and on new insights into learning that arise from computational models of intelligence [Papert 1973]. In the Logo laboratory, children are exposed to computers and computational concepts as a way of understanding and improving their own efforts to learn and to solve problems. By programming the computer to draw pictures, to play music, to simulate physical or biological processes, and to accomplish other substantive projects, students are introduced to important ideas in an active and concrete way.

The research proposed here benefits from the experience of the Logo Project, but represents a significant new line of research in several ways:

- \* **Theory:** Computer Coaching research requires the construction of *formal* theories of problem solving and of teaching, since these theories must serve as the basis of implemented modules in the Coach. The Logo group is primarily interested in the development of *informal* theories of problem solving, sufficient to guide a human teacher but not precise enough to serve as the basis of a tutorial program.
- \* **Experiments:** The Computer Coach allows tightly controlled experiments using the coaching system as a computational laboratory in which the modelling and tutorial components can be systematically varied. This is a new kind of experimental paradigm, not previously undertaken at Logo.
- \* **Tools:** The Logo project has focussed its attention on the development of computer languages and physical devices. This proposal is concerned instead with the incorporation of a tutoring component in the computer. Computer coaching adds a new kind of tool to the Logo environment: a cognitive advisor for the student.



An important caveat should be mentioned here. A major virtue of the Logo Project is that it provides environments where the student has enormous freedom. This is achieved by providing the child with a general purpose computer language and powerful peripherals. Traditional CAI, on the other hand, has often meant highly restricted environments for the students: only stereotyped replies were allowed or understood. We believe the computer game environment will provide sufficient freedom and opportunity for action that the player will not be unreasonably restricted, while the availability of a computer coach can be used to provide advice about underlying intellectual skills that the player can profitably use both in the game and in general. But we must be cautious not to fall into the trap of achieving artificial success by reducing the student's options to an intellectually uninteresting set. We can avoid this trap by allowing the computer game/coach environment to naturally grow into the full computer programming environment of Logo. This can be accomplished by allowing the student to design his own computer games (for which there would be no computer coaches) after having mastered the intellectual skills of those games for which tutors exist.

## 6. Critique

A number of potential objections to this research may have occurred to the reader. In this chapter, we reply to the more common reactions.

*The project is too ambitious. In the past, some proponents of CAI have overstated the potential of computers for education. Is this proposal a similar overstatement?*

CAI of the sixties attempted ambitious projects with limited hardware and software technology. Given the constraints on machine time and memory existing then, a project of the kind outlined here would have been impossible. But hardware is no longer a serious limiting factor: its costs are dropping drastically. A computational theory of tutoring, of modelling, of simplification does not yet exist. But there has been sufficient success in AI, in information processing psychology and in computational linguistics to make this research feasible. We believe two years will be sufficient to demonstrate the promise of this line of development.

There is another difference with traditional CAI. The coaching paradigm emphasizes that the learner as player is in control. It is not our intention to use the computer to return to a rigid format of "programmed instruction".

Finally, this research addresses fundamental questions of education. It provides a testbed for alternative theories of simplification, of explanation, of student modelling and human-oriented expertise. Traditional CAI did not focus on these fundamental issues.

*Skills acquired in games do not transfer!*

It is true that simply playing a game does not guarantee the acquisition of transferable skills. Indeed, repetitive play does not even guarantee perfect mastery of the game itself. It is for this reason that we believe Computer

Coaches have an important role to play. Their responsibility will be to gradually introduce the student to important concepts he has not discovered for himself.

Of course, not every game is a suitable arena for learning general intellectual skills. But we believe it is possible to design such games. Furthermore, by having families of games and a single Coach, the Coach itself could emphasize to the student the underlying regularities.

The danger that a fact learned in one context may not be applied in another is always present. We do not argue for Computer Coaches as the sole educational instrument. But appropriate games can supplement traditional presentation of mathematical or scientific material when supplemented by effective coaching.

We plan careful experimentation regarding this transfer issue. If Dewey was right that people learn by doing, we can expect positive results.

*Artificial intelligence research has not matured sufficiently to produce the level of performance demanded by this application.*

AI programs can already play games well, and perform expertly in certain arenas. The computer games discussed here are not more complex. What is more complex is communicating this expertise to a student/player. The bulk of this proposal addresses this issue by means of modelling, simplification rules, linguistic devices for concise discourse, and multiple explanation strategies. Whether these prove sufficient requires evaluation. But we feel this research lies on the critical path to taking educational advantage of the ongoing explosion of computer technology into our culture.

### 7. Conclusions

For many, the "athletics" model of learning -- games, teams, coaches, competition, skill -- is highly motivating and exciting. Traditionally only "physical" sports are taught in this way, with the more "serious" intellectual disciplines relegated to the classroom. Technology is making possible a new kind of sport: that of the computer-based intellectual game. Hence, the possibility exists for teaching the intellectual skills these games involve by means of the "athletics" model.

Computer games will be widespread within five years. The same experience we have seen with calculators with costs dropping to an insignificant level is about to recur for these TV games. Citizens of all ages will be playing and enjoying the sport they provide. Since these skills often involve basic mathematical and scientific knowledge, the player is acquiring an important kind of education in learning to play the game.

This essay has proposed a research program to investigate the design of computer coaches to facilitate the acquisition of intellectual skills exercised in these games. These coaches are far more complex programs than the games themselves, and it will take longer before every citizen can have his own coach. However, it would be reasonable for schools to provide such coaches, before they are sufficiently inexpensive for the home market. Indeed, such coaches might be an exciting drawing card for many students who are otherwise "turned off" by school. We would expect coaches to be affordable for schools by the early 1980's and for the home by the mid 1980's.

This research also has an important theoretical value. The design of a computer coach raises many questions central to psychology, to linguistics, to education and to artificial intelligence. By providing a common research focus, this project offers the possibility of a dynamic synergism between these fields.



Ultimately, it is our belief that applying the computational paradigm, as it has been developed in AI, linguistics and psychology, to education will contribute to a more powerful science of learning and of teaching.

### 8. Bibliography

- Barr, A., M. Beard and R. C. Atkinson, The Computer as a Tutorial Laboratory: The Stanford BIP Project, Technical Report No. 260, Psychology and Education Series, Institute for Mathematical Studies in the Social Sciences, August, 1975.
- Brown, J.S., "Uses of Artificial Intelligence and Advanced Computer Technology in Education", to appear in Proceedings of a Conference on the Impact of Computers on Education, A Ten Year Forecast, Washington, D.C., forthcoming 1976.
- Brown, J.S. & R. Burton, "Multiple Representations of Knowledge for Tutorial Reasoning", in D. Bobrow & A. Collins (Eds.), Representations and Understanding Studies in Cognitive Science, New York: Academic Press, 1975.
- Brown, J.S., R. Burton, & A. Bell, "SOPHIE: A Step Toward Creating a Reactive Learning Environment", International Journal of Man-Machine Studies, Vol. 7, 1975, pp. 675-696.
- Brown, J.S., R. Burton, & F. Zdybel, "A Model-Driven Question-Answering System for Mixed-Initiative Computer-Assisted Instruction", IEEE Transactions on Systems, Man, and Cybernetics, Vol. SMC-3, No. 3, May 1973, pp. 248-257.
- Brown, J.S., R. Rubinstein, & R. Burton, Reactive Learning Environment for Computer Assisted Electronics Instruction, Report 3314, Bolt, Beranek and Newman, Inc., Cambridge, MA, October, 1976.
- Buchanan, B., G. Sutherland and E.A. Feigenbaum, "Heuristic Dendral", in B. Meltzer and D. Michie (eds.), Machine Intelligence 4, New York, American Elsevier, 1969.
- Buchanan, B. et al, "Heuristic Theory Formation: Data Interpretation and Rule Formation", in B. Meltzer and D. Michie (eds.), Machine Intelligence 7, New York, American Elsevier, 1972.
- Burton, R, Semantic Grammar: An Engineering Technique for Constructing Natural Language Understanding Systems, BBN Report No. 3453, Bolt, Beranek and Newman, Inc., Cambridge, MA, December 1976.
- Burton, R. & J. S. Brown, "A Tutoring and Student Modelling Paradigm for Gaming Environments", in R. Coleman and P. Lorton, Jr. (Eds.), Computer Science and Education, (Advance Proceedings of the Association for Computer Machinery Special Interest Groups on Computer Science Education and

Computer Uses in Education Joint Symposium, Anaheim, CA), SIGCSE Bulletin, Vol. 8, No. 1, Feb. 1976, pp. 236-246.

Carbonell, J., "AI in CAI: An Artificial-Intelligence Approach to Computer-Assisted Instruction", IEEE Transactions on Man-Machine Systems, Vol. MMS-11, No. 4, December 1970.

Carbonell, J., Mixed-Initiative Man-Computer Instructional Dialogues, MIT Ph.D. Thesis, June 1970.

Carbonell, J. & A. Collins, "Natural Semantics in Artificial Intelligence", in Proceedings of the Third International Joint Conference on Artificial Intelligence, Stanford University, 1973, pp. 344-351.

Collins, A., Comparison of Two Teaching Strategies in Computer-Assisted Instruction, BBN Report No. 2885, Bolt, Beranek & Newman, Inc., Cambridge, MA, September 1974.

Collins, A., "Education and Understanding", in D. Klahr (Ed.), Cognition and Instruction, Hillsdale, NJ: Erlbaum Associates, 1975, pp. 287-289.

Collins, A., Processes in Acquiring Knowledge, BBN Report No. 3231, Bolt, Beranek & Newman, Inc., Cambridge, MA, January, 1976.

Collins, A., M.J. Adams, & R.W. Pew, The Effectiveness of an Interactive Map Display in Tutoring Geography, BBN Report No. 3346, Bolt, Beranek & Newman, Inc., Cambridge, MA, August, 1976.

Collins, A. & J. Carbonell, Semantic Inferential Processing by Computer, Bolt, Beranek & Newman, Inc., Cambridge, MA.

Collins, A. & M. Grignetti, Intelligent CAI, BBN Report No. 3181, Bolt, Beranek & Newman, Inc., Cambridge, MA, 1975.

Collins, A., J. Passafiume, L. Gould, & J. Carbonell, Improving Interactive Capabilities in Computer-Assisted Instruction, BBN Report No. 2631, Bolt, Beranek & Newman, Inc., Cambridge, MA, 1973.

Collins, A., E. Warnock, N. Aiello, & M. Miller, "Reasoning from Incomplete Knowledge", in D. Bobrow and A. Collins (Eds.), Representation and Understanding Studies in Cognitive Science, New York: Academic Press, 1975.

Collins, A., E. Warnock, & J. Passafiume, "Analysis and Synthesis of Tutorial Dialogues", in G. Bower (Ed.), The Psychology of Learning and Motivation, Vol. 9, New York: Academic Press, 1975.

Davis, R. Applications of Meta Level Knowledge to the Construction, Maintenance and Use of Large Knowledge Bases, Stanford Artificial Intelligence Laboratory, Memo AIM-283, July 1976.

Edelson, G.D., An Artificial Intelligence Analysis of Algebra Problem Solving, MIT B.S. Thesis, May 1976.

Goldberg, A. & P. Suppes, "A Computer-Assisted Instruction Program for Exercises on Finding Axioms",  
Educational Studies in Mathematics, 1972, 4, 429-449.

Goldstein, I. and E. Grimson, Annotated Production Systems, MIT AI Memo 389, in press.

Goldstein, I.P. and M. Miller, AI Based Learning Environments, MIT Artificial Intelligence Laboratory Memo  
384, Dec. 1976a.

Goldstein, I.P. and M. Miller, Structured Planning and Debugging: A Linguistic Theory of Design, MIT  
Artificial Intelligence Laboratory Memo 387, Dec. 1976b.

Greenblatt, R. "The Greenblatt Chess Program", Proc. FJCC (1967), pp. 801-810.

Greenblatt, Richard. The Lisp Machine. Working Paper 79, MIT Artificial Intelligence Laboratory,  
November 1974.

Grignetti, M., C. Hausmann, & L. Gould, "An 'Intelligent' On-Line Assistant and Tutor -- NLS-SCHOLAR",  
in Proceedings of the National Computer Conference, San Diego, 1975, pp. 775-781.

Grignetti, M. & E. Warnock, Mixed-Initiative Information System for Computer-Aided Training and  
Decision Making, Electronic Systems Division TR-73-290, September 1973.

Herot, C. F., Using Context in Sketch Recognition, Master of Science thesis, Electrical Engineering, M.I.T.,  
1974.

HUMRO, Proceedings of a Conference on the Impact of Computers on Education, A Ten Year Forecast,  
Washington, D.C., forthcoming 1976.

Marcus, M., "A Design for a Parser for English", in Proceedings of the Annual Conference of the Association  
for Computing Machinery, Houston, Texas, October, 1976.

McDonald, D., "A Preliminary Report on a Program for Generating Natural Language", in Proceedings of  
International Joint Conference on Artificial Intelligence, Tbilisi, USSR, August, 1976.

Miller, M. and I. Goldstein, Overview of a Linguistic Theory of Design, MIT Artificial Intelligence  
Laboratory Memo 383, December, 1976a.

Miller, M. and I. Goldstein, Parsing Protocols with Problem Solving Grammars, MIT Artificial Intelligence  
Laboratory Memo 385, Dec. 1976b.

Miller, M. and I.P. Goldstein, SPADE: A Grammar Based Editor for Planning and Debugging Programs,



MIT Artificial Intelligence Memo 386, Dec. 1976c.

Miller, M. and I. Goldstein, PAZATN: A Linguistic Approach to Automatic Analysis of Elementary Programming Protocols, MIT Artificial Intelligence Laboratory Memo 388, Dec. 1976d.

Negroponte, N., L. Groisser, and J. Taggart, Hunac, An Experiment in Sketch Recognition, Internal Memorandum, Architecture Machine Group, MIT, 1971.

Negroponte, N., Sketching A Computational Paradigm for Personalized Searching, Working Paper, Architecture Machine Group, MIT, 1975.

Papert, S., The Uses of Technology to Enhance Education, Proposal to the National Science Foundation available as MIT Artificial Intelligence Memo 298, June 1973.

Papert, S., An Evaluative Study of Modern Technology in Education, Proposal to the National Science Foundation, July 1976.

Purcell, S., Understanding Hand-Printed Algebra for Computer Tutoring, forthcoming Master's thesis, MIT, January 1977.

Resnick, C., Doctoral Dissertation, University of Illinois, 1975.

Rumelhart, D. E., "Notes on a Schema for Stories", in D. Bobrow and A. Collins (eds.), Representation and Understanding, New York: Academic Press, 1975.

Samuel, A. L., "Some Studies in Machine Learning Using the Game of Checkers", in E. Feigenbaum and J. Feldman (eds.), Computers and Thought, New York, McGraw Hill, 1963.

Shortliffe, E., Mycin - A Rule-Based Computer Program for Advising Physicians Regarding Antimicrobial Therapy Selection, Stanford AI Laboratory, AIM 251, October 1974.

Simmons, R., "Semantic Networks: Their Computation and Use for Understanding English Sentences", in R. Schank & K. Colby (eds.), Computer Models of Thought and Language, San Francisco: W. H. Freeman, 1973.

Smith, R. L., H. Graves, L.H. Blaine, & V.G. Marinov, "Computer-Assisted Axiomatic Mathematics: Informal Rigor," in O. Lecareme & R. Lewis (eds.), Computers in Education Part I: IFIP, Amsterdam: North Holland, 1975.

Stansfield, J., B. Carr, & I. Goldstein, Wumpus Advisor I: A First Implementation of a Program that Tutors Logical and Probabilistic Reasoning Skills, MIT Artificial Intelligence Laboratory Memo No. 381, Sept.

1976.

Suppes, P., "Computer Assisted Instruction at Stanford", in Man and Computer, (Proceedings of the International Conference, Bordeaux, 1970), Basel: Karger, 1972.

Winograd, T., Understanding Natural Language, New York: Academic Press, 1972.

Woods, W. et al., The Lunar Sciences Natural Language Understanding System, BBN Report 2378, Bolt, Beranek & Newman, Inc., Cambridge, MA, 1972.

Yob, G., "Hunt the Wumpus", Creative Computing, Sep-Oct, 1975, pp. 51-54.